

# Virtual Pooling Approximation Using Longest Path Network Optimization

by

Kevin E. Schell

B.S. Mechanical Engineering, Rice University, 2011

Submitted to the MIT Sloan School of Management and Department of Mechanical Engineering  
in partial fulfillment of the requirements of the degrees of

Master of Science in Mechanical Engineering  
and

Master of Business Administration

in conjunction with the Leaders for Global Operations Program at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2018

© 2018 Kevin E. Schell, All rights reserved

The author hereby grants to MIT permission to reproduce and to distribute publicly paper and  
electronic copies of this thesis document in whole or in part in any medium now known or  
hereafter created.

Author ..... **Signature redacted**  
MIT Sloan School of Management, Department of Mechanical Engineering

May 11, 2018

Certified by ..... **Signature redacted**  
Nikos Trichakis, Thesis Supervisor  
Zenon Zannetos (1955) Career Development Professor  
Assistant Professor of Operations Management

Certified by ..... **Signature redacted**  
Warren Seering, Thesis Supervisor  
Weber-Shaughness Professor of Mechanical Engineering

Approved by ..... **Signature redacted**  
Maūra Herson  
Director of MIT Sloan MBA Program, MIT Sloan School of Management

Approved by ..... **Signature redacted**  
Rohan Abeyaratne  
Chair, Mechanical Engineering Graduate Program Committee





# Virtual Pooling Approximation Using Longest Path Network Optimization

by

Kevin E. Schell

Submitted to the MIT Sloan School of Management and the MIT Department of Mechanical Engineering on May 11, 2018 in Partial Fulfillment of the Requirements for the Degrees of Master of Business Administration and Master of Science in Mechanical Engineering

## Abstract

This thesis proposes a network swapping mechanism to approximate a virtual finished goods inventory pool under the constraint of fragmented commercial channel ownership. Globally, Caterpillar sells its product through a network of independently owned dealers which own their equipment inventory. These dealers are selling a product which has significant configuration complexity, a high coefficient of variation in demand, and long factory lead times – these three factors conspire to create a situation in which dealers must either maintain high levels of on-hand inventory or sacrifice customer service level.

Virtual pooling can be a powerful technique for reducing held inventory and improving customer service level. In such a system, physical inventory is maintained in multiple locations, but a strong transshipment network allows inventory to be continuously rebalanced and effectively managed as a single pool. At Caterpillar, however, this is constrained by the fact that dealers are unwilling to unilaterally give up equipment that they own and which represents a potential sale, even if that equipment could be more effectively used by another dealer.

This thesis proposes that a robust dealer swapping network that allows for multilateral swaps can generate universally beneficial inventory movement which lowers inventory and accelerates sales across the network. The mathematical formulation of this problem involves solving a longest path problem over a suitably defined network. In order to demonstrate the efficacy of this technique, a commercial network model was developed which allows for the simulation of multi-period equipment ordering, inventory management, and sales across a sample dealer network with and without network swapping implemented.

Baseline simulation results conducted for a single vehicle class (Medium Wheel Loaders) suggest that network swapping has the potential to reduce on-hand inventory by more than 12% and decrease customer back orders by more than 17%. The swapping mechanism yields an NPV uplift of approximately USD 3 to 4M to the dealer network.

This thesis concludes by proposing important extensions of the work conducted in this thesis to improve the practicality and financial impact of the proposed network swapping system.

Thesis Supervisor: Warren Seering

Title: Professor of Mechanical Engineering, MIT

Thesis Supervisor: Nikos Trichakis

Title: Professor of Management Science, MIT Sloan School of Management

*This page intentionally left blank.*

## Acknowledgements

I would like to acknowledge the support of my MIT academic advisors, Professor Warren Seering and Professor Nikos Trichakis, for their guidance in structuring and communicating the research described in this thesis. I would like to thank Glenn Bergevin at Caterpillar, who seamlessly managed the multiple hats of supervisor, company thesis advisor, and friend. I would further like to thank my many other Caterpillar colleagues in Aurora, including Dave Shields, Lisa Ashley, Kip Mushisky, Kyle Kalmoe, Morgan Lynn, and Sarah Zameda.

I'd like to thank the Leaders for Global Operations staff and my LGO classmates for their guidance, ideas, encouragement, and support through this process.

Finally, I would like to specially acknowledge Donald Rosenfield, who – in conjunction with Warren Seering – advised the practical portion of this research and provided critical guidance in structuring the work which culminated in this thesis.

# Table of Contents

Abstract.....	3
Acknowledgements.....	5
List of Figures .....	9
How to Read this Document.....	11
1. Introduction .....	11
1.1. Background .....	12
1.1.1. Caterpillar.....	12
1.1.2. Construction Industries.....	13
1.1.3. Medium Wheel Loaders: Product .....	14
1.1.4. Medium Wheel Loaders: Value Chain.....	15
1.1.5. Engineered Value Chain .....	17
1.1.6. Stakeholder Interviews .....	18
1.2. Problem Statement.....	19
1.3. Scope.....	20
2. Literature Review.....	21
2.1. Kidney Paired Donation .....	21
2.1.1. History of Kidney Paired Donation.....	21
2.1.2. Types of Kidney Paired Donation .....	22
2.1.3. Parallels with Caterpillar Inventory Management.....	24
2.2. Virtual Pooling.....	26
3. Hypothesis.....	28
3.1. Current State Finished Goods Inventory Management .....	28
3.2. Proposed State Finished Goods Inventory Management.....	29
4. Methodology.....	31
4.1. Preliminary Sales Analysis.....	31
4.1.1. Sales to End User Analysis.....	31
4.1.2. Demand Correlation.....	36
4.1.3. Virtual Pooling Opportunity .....	38
4.2. Network Simulation Development .....	41
4.2.1. Overall Network Swapping Model Structure .....	41
4.2.2. End User Demand Modeling .....	44
4.2.3. Inventory Management Modeling.....	50

4.2.4.	Swapping Identification .....	51
4.2.5.	Model Validation Cases.....	57
4.3.	Dealer Syndication .....	59
5.	Simulation Results.....	62
5.1.	Baseline Network Scope .....	62
5.2.	Baseline Results .....	63
5.2.1.	Swap Characteristics .....	63
5.2.2.	Inventory Impact.....	65
5.2.3.	Backorder and Availability Impact .....	68
5.2.4.	Sales Acceleration .....	72
5.2.5.	Network Financial Impact .....	73
5.3.	Network Parameter Sensitivity .....	76
5.3.1.	Network Scale Effects .....	76
5.3.2.	Factory Lead Time Effects .....	79
5.3.3.	Complexity Requirement Effects .....	81
6.	Implementation Pathway .....	84
6.1.	Swapping Pilot.....	84
6.1.1.	Pilot Scope.....	84
6.1.2.	Pilot Logistics.....	85
6.1.3.	Pilot Process .....	85
6.1.4.	Pilot Organizational Ownership .....	86
6.2.	Durable System Development .....	87
6.2.1.	Inventory Monitoring System .....	87
6.2.2.	Bid/Ask Posting System.....	87
6.2.3.	Matching System.....	87
6.2.4.	Freight System .....	88
6.2.5.	Financial System.....	88
6.2.6.	Legal Framework.....	89
6.3.	Durable Process Development.....	89
7.	Conclusions and Extensions .....	90
7.1.	Conclusions .....	90
7.2.	Proposed Project Extensions .....	91
7.2.1.	Network Extension.....	91

7.2.2.	Alternative Swap Identification Techniques .....	92
7.2.3.	Complementary Initiative Impact .....	92
7.2.4.	Multi-echelon Problem .....	92
7.2.5.	Logistics Network Problem .....	93
References	.....	94
Appendix	.....	96
Executive Review Slides	.....	96
EVC Initiative Idea List	.....	103
Service Level Sample Analysis	.....	105
Dealer District Map	.....	106
Row-Wise Product Description	.....	107
Baseline Demand Scenarios, Gross Description	.....	108

## List of Figures

Figure 1: Example product families across Caterpillar segments .....	14
Figure 2: Table of MWL sales models sold in ADSD-N .....	15
Figure 3: Simplified finished goods value chain .....	17
Figure 4: KPD swap method summary .....	23
Figure 5: Analogous components of KPD and proposed swapping mechanism.....	25
Figure 6: Histogram of observed CV at a month / dealer / sales model level .....	33
Figure 7: Example Sales to End User series.....	34
Figure 8: Example simple service level model results.....	35
Figure 9: Dealer to dealer demand correlation summary .....	37
Figure 10: Coefficient of variation vs. mean monthly sales, example sales model .....	38
Figure 11: Preliminary virtual pooling model results, example model and district.....	40
Figure 12: Overall simulation model architecture .....	42
Figure 13: Simulated period commercial process order .....	43
Figure 14: Demand histograms with fitted Poisson distributions .....	46
Figure 15: Example empirical demand cumulative distribution function.....	47
Figure 16: Generating a synthetic demand series from the empirical CDF .....	48
Figure 17: Example configuration compression; five unique configurations are reduced to three by identifying three 'significant' configuration parameters .....	49
Figure 18: Illustrative swapping optimization scenario .....	51
Figure 19: Illustrative swapping network.....	52
Figure 20: The full set of network edges, A, and the durable transportation constraints, B .....	53
Figure 21: Matrix C, representing viable inventory/backorder pairings.....	54
Figure 22: Illustrative feasible swap quiver and directed graph.....	55
Figure 23: Illustrative reduction of matrix D (feasible swap quiver arcs) to matrix E (feasible swap directed graph arcs) .....	55
Figure 24: Model validation summary .....	58
Figure 25: Illustrative single period swap map .....	64
Figure 26: Overall swap characteristics, Baseline scenario .....	64
Figure 27: Total swaps classified by edge distance, Baseline scenario.....	65
Figure 28: Average on-hand inventory, Baseline scenario (Note: the average no-swap inventory has been normalized to 1,000 units to disguise data) .....	66
Figure 29: Single run on-hand inventory, swap vs. no-swap .....	67
Figure 30: Histogram of dealer count versus inventory reduction level .....	68

Figure 31: Average customer backorders, Baseline scenario .....	69
Figure 32: Single run backorders, swap vs. no-swap .....	70
Figure 33: Normalized backorder reduction vs. normalized no-swap backorders, Baseline scenario .....	70
Figure 34: Customer perceived availability, Baseline scenario.....	71
Figure 35: Single run cumulative sales advantage, swap vs. no-swap.....	72
Figure 36: Cumulative sales advantage, baseline scenario.....	73
Figure 37: Network swapping NPV, Baseline scenario .....	74
Figure 38: Single run cumulative, undiscounted incremental cash flows from network swapping.....	75
Figure 39: Single run Net Present Value of network swapping .....	75
Figure 40: Histogram of dealer level NPV improvement .....	76
Figure 41: Swap classification, baseline vs. intra-district scenario .....	77
Figure 42: NPV comparison, base case v. intra-district case .....	78
Figure 43: Undiscounted cash flow breakdown, intra-district scenario .....	79
Figure 44: Average monthly swap identification across various factory lead times .....	80
Figure 45: Network NPV impact across various factory lead times.....	80
Figure 46: Configuration parameters considered in 3-parameter, 6-parameter, and 9-parameter cases	81
Figure 47: Network swapping NPV across various configuration specificities .....	82
Figure 48: Undiscounted cash flows by source across different levels of configuration specificity .....	82
Figure 49: Summarized impact of configuration specificity on operating indicators.....	83

## How to Read this Document

This document is intended for multiple audiences, and the subsets of the following content are more relevant for certain audiences.

Individuals with an academic interest in the content described in this thesis – virtual pooling, network optimization, and complex inventory management – are encouraged to read this document in its entirety with the exception of Section 7, the Implementation Pathway, which is specifically targeted towards the initiative’s owners within Caterpillar. Each of the other sections provides important context, analysis, or discussion of the concepts that were explored in this thesis and the environment in which the described concepts may be applicable.

Individuals in general industry are encouraged to read the Abstract and focus on sections 1, 4, 6, and 8, which provide a description of the concepts set forth in this document as well as the potential impact that can be realized by implementing a network swapping system to approximate a virtual pool. Additionally, Appendix: Executive Review Slides provide a simple motivational description of the work set forth in this thesis.

Finally, individuals within Caterpillar who are seeking to further the body of work outlined in this document are encouraged to read the entire document with a particular focus on sections 7 and 8, which outline the proposed implementation pathway and further development horizons for the work described in this thesis.

### 1. Introduction

How does a firm minimize the resources required to ensure that it has the right product in the right place at the right time to make a sale?

This question is the nucleus of value chain design. It’s also deceptively complex – solving this puzzle has been a foundation of operations research for decades, and the diversity that can be observed in real-world value chains is a testament to the fact that no single optimum answer exists. However, history suggests that the potential impact of finding ways to differentiate and improve a value chain is compelling. In the late nineteenth century, John D. Rockefeller founded and consolidated one of the largest companies the world has ever seen – the Standard Oil Company – by developing a rail and pipeline logistics network that created an insurmountable cost advantage over its competition’s fragmented value chain (Chernow, 1997). Seventy years later, one of the largest supply chains ever conceived delivered over one million soldiers, 560 thousand tons of supplies, and 170 thousand vehicles to Normandy during the D-Day invasion by using innovations such as artificial harbors and the PLUTO

pipeline [2]. At the turn of the millennium, Amazon arose from a sea of also-ran .com companies to become the largest retailer in the world; it did so by developing a value chain with unparalleled speed, breadth, and cost.

On an immediate scale, value chain improvement is about making one more sale or delivering a product one day earlier. This thesis examines the finished goods value chain of a single company, Caterpillar. It proposes a novel swapping mechanism which would provide Caterpillar's dealers an additional tool for making existing sales more quickly and for capturing incremental sales.

## 1.1. Background

The network swapping concept discussed in this thesis was developed in conjunction with Caterpillar. The proposed solution was developed to address challenges specific to the context described in this section.

### 1.1.1. Caterpillar

Caterpillar traces its heritage to 1904, when Benjamin Holt first built and tested a track-type tractor designed to work in agricultural applications where wheeled tractors would sink into soft soil. In 1925, Holt Manufacturing Company merged with a competitor, C.L. Best, and the name of the merged company was changed to the Caterpillar Tractor Company to reflect the mode of locomotion for its vehicles [3].

Today, Caterpillar, Inc. manufactures heavy equipment for the construction, resource, and energy industries. Broadly, Caterpillar classifies its businesses into three segments: Construction Industries, Resource Industries, and Energy & Transportation. The Construction Industries segment encompasses equipment that primarily operates in infrastructure, forestry, and building construction operations (e.g. excavators, medium bulldozers, medium wheel loaders, motor graders, etc.). This segment accounted for approximately 44% of Caterpillar's core business revenue in 2016. The Resource Industries segment is composed of large equipment that primarily operates in extractive industries (e.g. mining shovels, off-highway trucks, large wheel loaders, and large bulldozers) and accounted for 16% of 2016 core business revenue. The energy & transportation segment primarily consists of reciprocating engines, gas turbines, and diesel-electric locomotives. This segment accounted for 40% of 2016 core business revenue [4].

Caterpillar competes globally against companies such as Komatsu, Volvo, Hitachi, John Deere, and XCMG. In most of the segments where it competes, Caterpillar is engineered and marketed as a premium product; its machines are designed to operate for more cycles in heavier applications with higher efficiencies than competing products, and its equipment frames are engineered to support several rebuilds (or "product lifetimes"). Despite selling at a premium relative to its competitors,

Caterpillar still commands a leading market share; Equipment World estimated that Caterpillar held an 18.1% market share in 2015 (Komatsu, in second, held a 10.5% market share) [5].

Caterpillar's dealer network is an integral component of its commercial strategy. With relatively few exceptions (e.g., corporate accounts, Solar Turbines), Caterpillar sells the equipment that it produces to a network of independently owned dealers around the globe. These dealers cover mutually exclusive geographic territories, managing all customer relationships, sales, and maintenance within that geography. Many of Caterpillar's dealers are large corporate entities unto themselves, and several are publicly listed. Publicly available press clippings highlight the fact that this dealer network has, at times, been a mixed blessing. While Caterpillar credits their dealers with building deeper relationships and customer understanding that it would be able to itself, it also has recognized instances in which the decision to abdicate control over the final link in its value chain has prevented an agile response to changing market dynamics. [6], [7]

#### 1.1.2. Construction Industries

The various industries in which Caterpillar competes behave quite differently. At one end of the spectrum are Resource Industries products. In the extreme case of draglines, these machines are capital investments whose sale is more similar to that of a building than most machinery; customers would contract up front to build a machine which would be engineered and built to meet their specific requirements. Although Resource Industries presents a significant long-term capacity planning challenge (due to the cyclicity of underlying industries), machines are built to order on a short-term time frame. Consequently, the sales planning cycle is simplified, and neither Caterpillar nor its dealers carry significant inventory.

Within Construction Industries, Caterpillar internally further divides its product into two separate categories: Building and Construction Product (BCP) and general Construction Industries (CI). BCP products are small (typically less than 15 tons), relatively low cost, and high volume machines. A prototypical BCP machine is the backhoe loader. Configuration options for BCP machines are relatively limited, and customers often make same day purchases from dealers' inventory.

General CI machines, which this thesis focuses on, exist in the intermediate space between Resource Industries and BCP machines. Many iconic Caterpillar machines – excavators, medium bulldozers, medium wheel loaders, and motor graders – fall under the aegis of general CI. General CI machines represent significant investments for many customers; as a consequence, customers expect to have a machine that is specifically customized to meet their application. Furthermore, volumes are relatively low; a typical dealer may sell anywhere from one to ten machines of a given general CI model in a

month. Despite this, competitive pressure and customer requirements dictate that most machines in this size range are still sold from stock, and dealers carry significant inventory levels of general CI product. These factors make value chain management in the general CI segment particularly challenging and require close cooperation between Caterpillar and its dealers to ensure that product is in the right place at the right time to make a sale.



Figure 1: Example product families across Caterpillar segments

### 1.1.3. Medium Wheel Loaders: Product

The concepts discussed in this thesis are broadly applicable to Caterpillar’s general CI category, and may be extended to encompass BCP product, as well. However, the initial research and results outlined in this thesis were assembled using data for a single product family: Medium Wheel Loaders.

Medium Wheel Loaders (MWLs) are center-articulated wheeled tractors with a single front mounted work tool. MWLs are largely used in industries such as infrastructure, logging, aggregate handling, and

small-scale mining. MWLs are one of the highest volume products in the general CI category, and the model family has sufficient breadth and configuration complexity to function as a research proxy for the full range of general CI products. Caterpillar currently produces eight sales models of MWLs for the ADSD-N market, which are summarized in the table below.

Sales Model	Operating Weight	Net Engine Power	Commentary
950M	42,357 lb.	230 HP	CAT recently began selling a utility segment '950GC' which is considered a separate sales model
962M	44,591 lb.	250 HP	Shares a common frame with 950M
966M	51,176 lb.	276 HP	
972M	54,871 lb.	299 HP	Shares a common frame with 966M
980M	66,337 lb.	386 HP	
982M	78,402 lb.	389 HP	
966MXE	51,176 lb.	298 HP	Advanced drivetrain 966 (CVT); relatively low volume sales model in ADSD-N
972MXE	54,871 lb.	311 HP	Advanced drivetrain 972 (CVT); relatively low volume sales model in ADSD-N

Figure 2: Table of MWL sales models sold in ADSD-N

Configuration complexity is an important characteristic of the machines that Caterpillar sells. In essence, significant configuration complexity further fragments the demand signal, because major configuration variants cannot be substituted for each other. In the MWL product family, major configuration variants include dimensions such as boom (short-lift, standard-lift, high-lift, and logger), hydraulics (2-valve, 3-valve, or 4-valve arrangements which dictate work tool control), differential (manual locking vs. auto locking), and axle oil cooler (present or absent). MWLs have approximately 40 dimensions of configuration, but only a handful of these are considered major by customers and dealers.

#### 1.1.4. Medium Wheel Loaders: Value Chain

Figure 2 is a simplified conceptual diagram for the Medium Wheel Loaders value chain focused on orders and fulfillment. There are two regimes under which the finished goods value chain operates: free supply and allocation.

Under free supply, dealers are fully responsible for their own inventory planning. Dealers typically set target inventory levels based upon a desired *months* of typical sales volumes loosely based upon a forward coverage methodology. This forward coverage methodology is encouraged via a Caterpillar system known as DDVS, which – using on-hand stock, forecast sales, and factory lead time – provides dealers with a recommendation of how many machines to order each month.

Under allocation, Caterpillar sets a maximum factory lead time that it holds fixed. Caterpillar manages this lead time by rationing build slots across dealers based upon their sales and remaining inventory.

Allocation is an attempt to manage a behavioral reinforcing feedback mechanism that exists when factory lead times grow longer: dealers exhibit a tendency to over-order because of the perception that the next order will take even longer for them to receive. Once demand begins to recede, this behavior creates an overhang of excess inventory in the system. The allocation process is managed by the MWL S&OP team.

When submitting an order, a dealer may use one of three order *lanes*, which dictates the lead time required to receive a machine. Lane 3 is a fully customizable order – under free supply, dealers can submit Lane 3 orders at any time and can configure a machine according to any option available on the Caterpillar published price list (under allocation, dealers are told how many Lane 3 machines they can order). With the exception of a few low-running options, Lane 2 orders are also highly configurable; however, a dealer must commit to a program known as Base Order Management (BOM) in order to use Lane 2. Under this program, a dealer has an ongoing commitment to purchase a certain number of machines per month, and those machines will be delivered via Lane 2 with a fixed, guaranteed ‘lead time’ (Note: the metric that Caterpillar actually fixes is internally termed *availability*, and measures the time from order receipt to shipment; shipping time itself is not included) that is shorter than Lane 3. The theory of BOM is that it helps manage behaviors that contribute to so-called value chain bullwhip – a phenomenon in which variability in the demand signal increases as one progresses upstream in the value chain. [8] Lane 1 orders have limited configurability and are delivered from a relatively small amount of inventory that Caterpillar itself builds to stock and owns. These standardized machines are intended to provide surge capacity in the event that one or more dealers face unexpectedly high demand and are consequently intended to act as a ‘safety valve’ that dealers only use to pursue fast moving sales that cannot be met through their own inventory. As an approximation of the relative lead time for these three ‘lanes,’ a Lane 3 machine may have a three-month lead time, a Lane 2 machine may have a two-month lead time, and a Lane 1 machine may have a lead time of less than one month.

Once production is complete, the logistics of order fulfillment are managed by Caterpillar’s Global Supply Network Division (GSND). Within the Construction Industry tier of product, most final delivery is executed via Caterpillar’s Last Mile network of contracted flatbed trucks. The size of finished machines in this segment are such that shipments are typically unit of one; depending on the destination, larger CI machines may need to be partially disassembled for shipment.

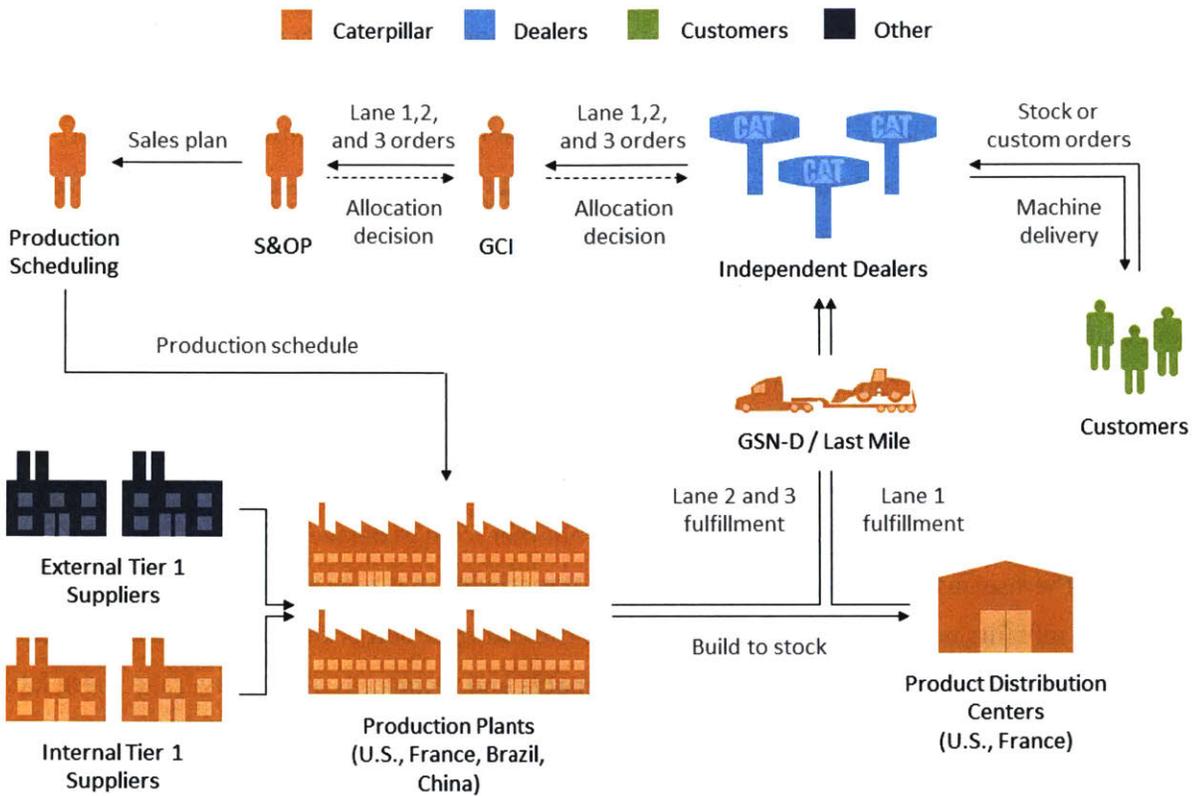


Figure 3: Simplified finished goods value chain

### 1.1.5. Engineered Value Chain

Caterpillar launched a corporate initiative in 2014 known as the Engineered Value Chain (EVC). The stated mission of EVC, according to Caterpillar, is the following:

*We must engineer our value chain, much like we would engineer our product. Starting with the customer requirements and designing the entire value chain based on those requirements back to the engineer's pen.*

Organizationally, EVC is owned centrally by the GSND group. However, the implementation of EVC is driven by product groups (e.g., Medium Wheel Loaders) themselves. Perspectives on the ultimate goal and efficacy of EVC vary within Caterpillar; centrally, work has been done which focuses on providing the *appropriate* value chain channel for each customer. Central voice of customer work suggests that some customers are relatively indifferent to product lead time and are willing to wait significantly longer to receive their machine if there is any associated cost savings. Conversely, some customers are very lead time sensitive and are willing to pay a premium or switch suppliers to receive a machine more quickly. This central perspective holds that Caterpillar should create different value chain channels in order to meet different customer needs.

In practice, product groups have translated EVC has an invocation to universally reduce value chain cost and lead time (these two goals are, obviously, sometimes in conflict, and a clear primacy does not exist). This interpretation has yielded initiatives such as Finish to Order (FTO), dealer kitting, and product containerization. FTO is the most widely adopted initiative to arise from the EVC program (Last Mile – Caterpillar’s terminal logistics network – is also now considered a part of EVC, but it predated EVC) and, in some cases, has become synonymous with EVC. FTO refers to the process of delaying certain aspects of product differentiation until after the machine has left the product line but before it is shipped to a customer. Within the MWL product group, FTO is primarily circumscribed to Lane 1 machines and occurs at the Product Distribution Centers, allowing the MWL group to offer a wider range of configurations for its Lane 1 machines. This theoretically allows dealers to pull on the accelerated Lane 1 delivery pathway for a wider range of fast moving deals. Dealer kitting is another form of delayed differentiation which allows dealers themselves to modify the configuration of machines that they already own. Caterpillar assembles and sells kits which dealers can use to revise the configuration of their machines after receipt. Although most of Caterpillar’s dealers are extremely capable and technically able to leverage these kits, significant adoption has been circumscribed to a handful of dealers. Finally, containerization is an effort to leverage the PDC rework capability developed for FTO in order to conduct final product assembly after shipment from an overseas production facility. The objective of this EVC initiative is to lower shipment costs.

#### 1.1.6. Stakeholder Interviews

Prior to focusing on a single aspect of the value chain, a series of approximately fifteen interviews were conducted with stakeholders in the Medium Wheel Loader product group. Additionally, Caterpillar facilitated an introduction to one of its highest volume MWL dealers, and a visit was conducted to interview additional stakeholders within the dealer itself.

The first major theme that arose during these interviews was the perspective that product configuration complexity is a major challenge for Caterpillar broadly and the MWL product group specifically. Internally, Caterpillar stakeholders spent significant amounts of time highlighting the time and work that is driven by configuration complexity: engineering must maintain a large volume of drawings for each machine, accounting is unable to create standard costs that translate across the configuration range, and the supply chain faces disruptions driven by tier 1 suppliers supplying the incorrect arrangement for highly configured sub-systems. Dealers, on the other hand, struggle with a price list that – for each sales model – exceeds thirty pages and can be difficult to interpret; consequently, many dealers set a few of their own standard arrangements (which, unfortunately, deviate from the standard arrangements of

other dealers) and default order those arrangements in order to avoid revisiting the price book regularly.

The second theme identified in these interviews is that Caterpillar and its dealers face significant agency problems driven by fragmented ownership. Although Caterpillar often pursues initiatives which are primarily intended to benefit its dealer network, those initiatives are inevitably more challenging to justify than initiatives which have a clear bottom line impact for Caterpillar itself. Unfortunately, many EVC initiatives fall into this category and face a higher effort gradient to achieve approval. Among dealers, the same agency problems exist, particularly with respect to sharing inventory. Informally, dealers will occasionally transfer equipment with each other to facilitate sales; however, this is typically done as the consequence of a standing personal relationship between sales operations managers of dealers within a region and often takes a *quid pro quo* format over time.

Finally, Caterpillar and its dealers both maintain the perspective that the right inventory – whether that inventory is parts, attachments, or machines – usually exists *somewhere* within the Caterpillar-dealer ecosystem, even if it is not in the correct place.

## 1.2. Problem Statement

Caterpillar's Medium Wheel Loader Group is characterized by a complex product, long lead times, and a fragmented commercial channel. Each MWL sales model has approximately 40 different dimensions along which it can be configured, resulting in millions of possible machine configurations. The complexity of the MWL product, combined with the scope of Caterpillar's supply chain results in lead times from Caterpillar to its dealers of several months. Finally, Caterpillar has approximately 50 dealers in North America alone, each of which owns and manages its own pool of finished goods inventory.

Together, these factors conspire to create high levels of finished goods inventory. In 2016, average MWL finished goods inventory exceeded 15 weeks of sales and \$250 MM in value. Anecdotally, these factors also suppress customer service level (defined as the percentage of customers which are served directly from inventory rather than having to wait for a machine) at the point of sale - data tracking among dealers is currently too inconsistent to corroborate this empirically.

Seventeen potential initiatives were identified which could be pursued to help address these challenges (see Appendix: EVC Initiative Idea List). Based upon the charge to develop an initiative which is orthogonal to existing Engineered Value Chain initiatives, this project was scoped to focus on addressing the challenge of fragmented ownership and management of finished goods inventory. **The objective of the research work documented in this thesis was to develop a swapping mechanism which would**

allow Caterpillar and its dealer network to approximate an inventory pool under the constraint of fragmented inventory ownership.

### 1.3. Scope

The scope of this project is circumscribed in several dimensions.

Geographically, this project will focus on the ADSD-N marketing division within Caterpillar. The ADSD-N marketing division is composed of 47 dealers covering the United States and 4 dealers within Canada. The rationale for this scope is three-fold. First, customer and dealer behavior and expectations are somewhat consistent across this region, so it is feasible to meaningfully obtain representative feedback from a subset of stakeholders. Second, the trans-shipment logistics within this region are feasible, both with respect to geographic continuity and customs constraints. Finally, ADSD-N is the largest region for Medium Wheel Loader sales globally, accounting for 35-40% of annual volume.

As discussed in the Introduction, this research also focuses specifically on Medium Wheel Loaders. This scope constraint was primarily driven by the accessibility of this product group. The concepts outlined in this thesis are applicable to the general Construction Industries product range, and the potential impact of including additional product groups is likely super-linear (as discussed in Results).

Finally, the results outlined in this thesis are focused on a single generation of product (specifically, M-series loaders). In the proposed swapping system, there are likely transitory effects associated with the cut-over from one generation of product to the next (e.g., the configuration options between two generations of product will be different). These transitory effects are not considered.

## 2. Literature Review

The network swapping mechanism proposed in this thesis was inspired the Kidney Paired Donation (KPD) system which was first proposed in 1986 and has been used in the United States to facilitate kidney donation matches since 2000. When applied to the inventory problem faced by Caterpillar, the concepts behind KPD provide a technique which allows Caterpillar to approximate a virtual inventory pool across its independently owned dealer network by using integer programming to identify inventory swaps.

### 2.1. Kidney Paired Donation

Kidneys are the most commonly donated organ, accounting for more than half of all successful organ donations in the United States. Despite this, the requirement for kidneys significantly exceeds the supply: the wait-list for kidney donation in the United States is currently nearly 100,000 patients, while 19,849 kidneys were successfully transplanted last year. [9] One of the major contributors to this mismatch is incompatibility: while a patient requiring an organ may have a willing donor, that donor's organ may not be biologically compatible with the recipient. In this case, Kidney Paired Donation (KPD) has arisen as a method for securing donations for patients whose willing donors are not a biological match.

#### 2.1.1. History of Kidney Paired Donation

KPD was first proposed by Dr. Felix Rapaport, a former surgeon and faculty member at NYU, who proposed a "living emotionally related international kidney donor exchange." [10][11][12] The concept set forth by Dr. Rapaport was first translated into reality in 1991 in South Korea when two living donor-recipient pairs 'swapped' kidneys in a simultaneous surgery; in order to realize biologically compatible transplants, each willing but incompatible donor donated a kidney to the opposite donor's relative as opposed to their own relative. [13] This first in the world program in South Korea resulted in 129 transplants over an 11-year period. In 2001, the South Korean network was formalized nationally across the country's transplant centers with the formation of the National Korean KPD program. [14]

Since 1991, KPD has expanded globally, with each region developing its own exchange networks. Within Europe, the first paired kidney swap was conducted in Switzerland in 1999, and the first formalized national program – coordinating patient pools across seven transplant centers – was established in the Netherlands in 2004. The United States (which faced higher legal barriers KPD as a consequence of an idiosyncrasy of the 1984 National Organ Transplant Act) followed, with its own first paired donation occurring at Rhode Island Hospital in 2000. [14] A nationally pooled system didn't arise until 2010 with the formation and pilot of the United Network for Organ Sharing (UNOS). [15][14]

In the early stages of KPD, paired donations were identified manually. In 2004, a significant advancement in KPD was realized when Alvin Roth et. al. formalized the concept of a centrally cleared, algorithm-based swap identification program. [16] This proposed clearinghouse ultimately became the New England Program for Kidney Exchange (NEPKE), which was the first computerized matching system for KPD. Between 2005 and 2010, this system resulted in 58 successful transplants. [14]

Since the success of NEPKE, algorithmically directed, centrally cleared systems have become the dominant method of KPD, and these systems have yielded swaps that would have been impossible to identify manually. In 2012, for example, the New York Times reported on a donation chain facilitated by the National Kidney Registry which consisted of thirty swaps across sixty patients. [17] These algorithmically arranged swapping networks have also dramatically improved the impact of individual donors; a 2013 study found that the average non-directed donor (i.e. a kidney donor willing to be matched to any potential recipient) resulted in nearly five successful transplants. [18]

#### 2.1.2. Types of Kidney Paired Donation

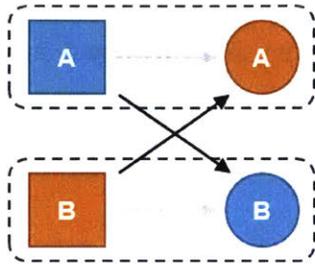
As KPD matching systems have matured, a number of strategies have emerged for increasing the density and quality of kidney swaps identified by the exchanges. Figure 4 summarizes several of the major swap mechanisms employed in KPD.

**Diagram of KPD Swap Techniques**

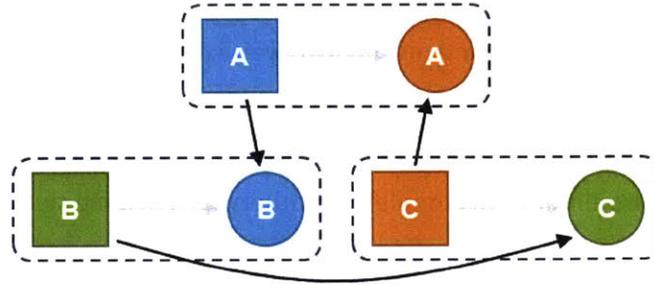
n = number of month / dealer / sales model atoms in CV range



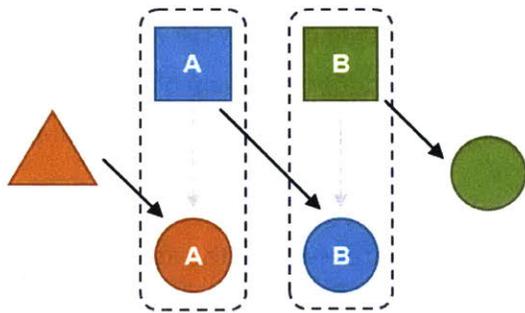
**a. 2-way exchange**



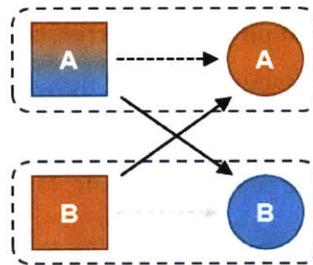
**b. 3-way exchange**



**c. Domino-paired donation**



**d. Compatible pair donation**



**e. NEAD chain**

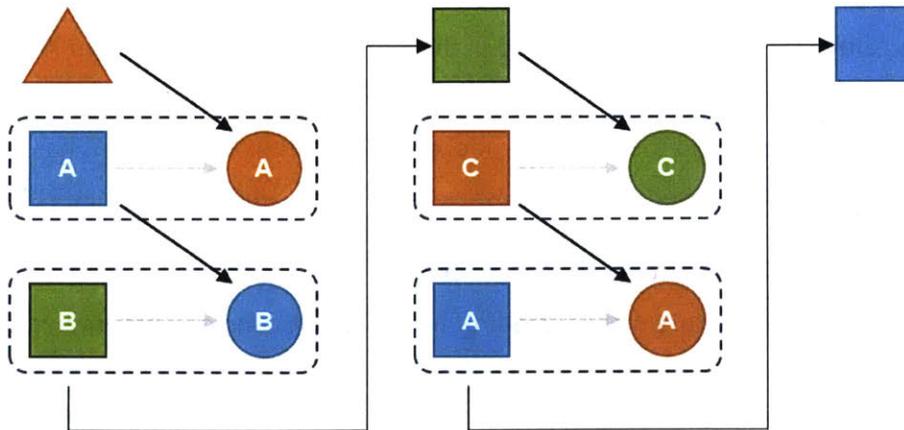


Figure 4: KPD swap method summary

The first kidney swap performed in South Korea in 1991 was done according to method of Figure 4a – two willing donors (spouses of the donor recipients) were not biological matches for their intended recipient but were matches for each other’s intended recipient. This allowed for a simultaneous swap in which each donor’s kidney was implanted in the other donor’s spouse. [14] A nuanced extension of this swapping mechanism is presented in Figure 4c in which a viable donor-recipient pair instead takes part in a matched swap. This mechanism can simultaneously provide a kidney for the pair which is otherwise unmatched while improving compatibility for the matched pair. [16]

The bilateral swapping mechanism was extended in 1995 to include three- and four-way swaps. [11] In the case of multi-way swaps (shown in Figure 4b), additional donor-recipient pairs are included in the swap pathway to either improve the match (i.e., reduce the risk of organ rejection for one or more of the recipients) or to complete an otherwise unviable swap chain. [17]

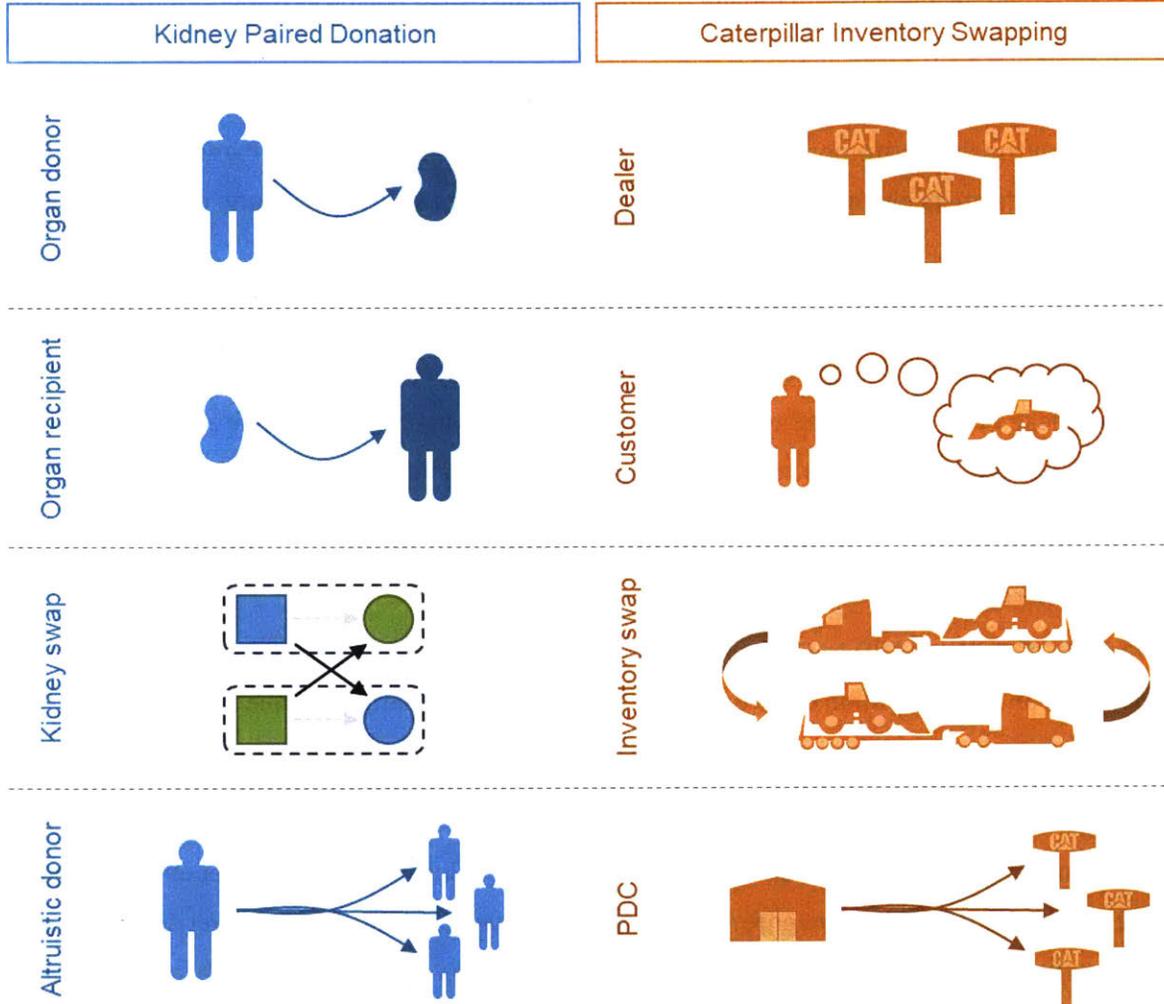
An important component of KPD networks are individuals known as altruistic, non-directed donors. These healthy individuals are willing to donate one of their two kidneys to a patient with whom they have no relationship. When applied to a kidney swapping program, these altruistic donors dramatically increase the number of potential matches by allowing acyclic swapping pathways in which a chain of swaps can be started anywhere in the network and the last willing donor in a chain does not have to be a match for the first recipient in the chain. This mechanism, which is sometimes termed a domino-paired donation can be seen in Figure 4c. [11][14]

One of the major challenges for KPD – particularly in the United States – is that the associated surgeries must be performed simultaneously. Since federal law prohibits contracting for human organs, simultaneous surgeries is the only way to completely ameliorate the risk of one of the donors in a chain from reneging once his or her targeted recipient has received a transplant. However, when the risk of a donor reneging is low, altruistic donors can be applied in the matching mechanism shown in Figure 4e. In this case, known as a non-simultaneous extended altruistic donor (NEAD) chain, the donor for the terminating pair in the chain remains in the matching pool even after his or her recipient has received an organ, and that terminating donor is used as the altruistic catalyst for a subsequent chain. [14]

### 2.1.3. Parallels with Caterpillar Inventory Management

Upon cursory examination, the inventory management of heavy construction equipment bears little resemblance to organ donation and KPD. However, parallels can be drawn between these two problems which suggest that the concepts underpinning KPD have the potential to significantly impact inventory management across Caterpillar’s dealer network. This extended analogy is summarized in Figure 5 below.

**Approximate analogies between KPD and Caterpillar swapping network**



*Figure 5: Analogous components of KPD and proposed swapping mechanism*

A non-compatible kidney donation pair occurs when a patient in need of a kidney has a willing donor whose inclination to donate is circumscribed to that specific patient (e.g., an individual that is willing to donate to an ill spouse). In the absence of KPD, the only recourse for the non-compatible recipient is to wait for an organ from a deceased organ donor. The Caterpillar counterpart for this phenomenon occurs when a customer arrives at the single Caterpillar dealer serving his or her geographic region and seeks to order a machine configuration that the dealer does not have in stock. When this occurs, the customer can choose to wait for the Caterpillar dealer to order and receive the desired machine configuration, or the customer can pursue competitive offerings. The application of the KPD model to Caterpillar’s dealer network opens up a third alternative.

In the stylized example above, the Caterpillar dealer is the non-compatible donor, while the customer is the non-compatible recipient; compatibility is defined by the dealer having a configured machine which meets the customer's purchase requirement. There is some likelihood that another dealer in Caterpillar's network (another 'donor') has a configured machine on hand which matches the requirements of the first customer. If that matching dealer also has customer demand which cannot be immediately met from its inventory, then the potential for a swap similar to the one outlined in Figure 4a arises.

This analogy can be extended through the inclusion of the Caterpillar owned inventory housed at its Product Distribution Centers. Since Caterpillar does not generally sell to customers directly, the inventory in these PDCs can be used to approximate the role of an altruistic non-directed donor – PDC machines can be sent to any dealer to initiate a swapping chain, and the PDC does not need to receive a machine from the terminal dealer in the swapping chain.

## 2.2. Virtual Pooling

The advantage of applying KPD to Caterpillar's inventory management problem derives from the fact that it provides an approximation of inventory pooling where fragmented physical pools *and* fragmented inventory ownership across dealers would otherwise prevent pooling.

In 1979, Gary Eppen of the University of Chicago published a paper which set forth one of the most important concepts in operations management: for  $n$  inventory locations facing independent, identically distributed demand, the level of inventory required to achieve a targeted service level declines according to the square root of  $n$  when the multi-site inventory is pooled into a single location. [19] Since this result was obtained, legion papers have extended the concept of inventory pooling. One important category of extensions examines the case of virtual pooling or lateral transshipment. In the case of virtual pooling, pooled inventory is not held in a single location but is transferred between 'virtually pooled' locations when needed according to a logistics network. The benefit of such virtual pools must weigh the tradeoff of increased logistics cost and decreased inventory cost. [20][21]

Two important dimensions along which virtual pooling methods can be classified are whether the virtual pooling is proactive or reactive and whether the method provides complete or partial pooling. Proactive systems redistribute system inventory according to a fixed cadence, while reactive systems react to stockouts or the threat of stockouts at one or more locations in the network. Proactive systems are generally more appropriate for networks where logistics costs represent a significant portion of product costs, and the volume of units shipped at any one point in time must be sufficient to justify the network's logistics cost. Reactive systems are appropriate for situations where logistics represent a

relatively small portion of a product's cost. Systems with complete pooling allow inventory locations to share their full inventory with any other inventory location in the system, while partial pooling systems create some restrictions in which an inventory location reserves some inventory to cover its own future demand. [21] The system discussed in this thesis is a form of reactive, partial pooling.

In 2014, Smirnov and Gerchak examined a form of virtual pooling which resembles the logistical constraints of KPD and the method proposed in this thesis. They examined a system where each inventory location in a lateral transshipment network was circumscribed to receiving inventory from a single antecedent node and sending inventory to a single subsequent node. For a three node system, their result showed that this 'unidirectional chaining' increased logistics costs by 19% over a completely pooled network in which each node could ship freely to the other two nodes. [20]

### 3. Hypothesis

The network swapping mechanism proposed in this thesis is intended to provide some of the benefits of true virtual inventory pooling within the constraints imposed by the organization of Caterpillar's independent dealer network. Specifically, it is hypothesized that the proposed swapping mechanism will do the following:

- Result in a **one-time reduction of on-hand inventory** that is less than what would be realized in the case of a true virtual pool but sufficient to justify implementation of the swapping network
- Result in a **one-time acceleration of sales to end customers**, the benefits of which will accrue to the dealers (which already own the inventory that is being 'pooled')
- Result in a revenue uplift – derived from sales acceleration – and inventory holding cost savings which exceeds the net present value of the logistics costs required to support the swapping network; in other words, the **proposed initiative will be NPV positive**
- **Reduce the perceived delivery lead time for a subset of customers**, which – in turn – will improve Caterpillar's market share among customers who are time sensitive; this represents an ongoing financial uplift incremental to that achieved through one-time optimization

Furthermore, it is hypothesized that certain sensitivities will exist within the proposed swapping network:

- The **benefits of the proposed swapping network will grow approximately linearly with the number of dealers** included in the network; the justification for the hypothesis of linearity is that the number of viable swap edges will grow according to the square of the nodes in the underlying network while the strict benefits of pooling (in the limiting case of uncorrelated demand signals) grow according to the square root of inventory volume
- The proposed swapping network will be **highly sensitive to the degree of configuration specificity** required by dealers; adding degrees of configuration specificity will reduce the likelihood that a matching machine exists within the network
- **Inclusion of the Caterpillar inventory at the Product Distribution Center as an altruistic donor will improve the benefit** realized in the proposed swapping network by allowing swapping vectors in addition to circles

#### 3.1. Current State Finished Goods Inventory Management

The relevant characteristics of the current state of finished goods inventory management at the dealer level can be characterized as follows:

- When formally sourcing inventory, dealers have **two primary delivery pathways** which they can utilize: they can order fully-customizable inventory from a Caterpillar factory with a relatively long lead time or semi-customizable inventory from a Product Distribution Center (PDC) with a relatively short lead time.
- The typical **review period for the standard ordering process is one-month** (driven by the DDVS cadence), although this review period is compressed when a dealer faces potential fast-moving deals.
- Dealers will **occasionally informally source inventory via ad hoc swaps**. These ad hoc swaps are sourced in two ways:
  - A dealer that needs a specific machine will consult the Machine Products Inventory Manager (MPIM) which offers limited visibility into other dealers' inventory (dealers can suppress the visibility of their inventory if they wish) to find a machine which matches their requirements. The asking dealer will then call or email the dealer which has the desired unit of inventory.
  - A dealer that needs a specific machine will send a blanket email to the entire dealer set within a given region.

In either case, the dealers will **directly negotiate the terms of a swap with each other and arrange the logistics of machine transport** themselves. This informal swap network is largely driven by personal relationships between dealers and the history of exchanges between specific dealers (i.e., one dealer may feel that they 'owe' another dealer because of a previous inventory swap).

### 3.2. Proposed State Finished Goods Inventory Management

The proposed state for finished goods inventory management seeks to (1) formalize the current ad hoc swapping network into a common process shared by all dealers and (2) lessen constraints on dealer-to-dealer swaps by allowing multi-lateral swapping. The relevant characteristics of the proposed system are the following:

- Dealers will have **three possible avenues through which they can formally source inventory**. The factory sourcing process will remain unchanged. Dealers will also continue to have access to the PDC through the Lane 1 ordering process; however, **PDC inventory will be preferentially used to serve as a starting node for open-ended swapping chains in the proposed process** (this role is similar to that of an 'altruistic donor' in Kidney Paired Donation; henceforth we will use this term to describe the role the PDC plays in the swapping network). Finally, dealers will be able to regularly submit swap requests to source short lead-time inventory from other dealers.

- The review period for the standard ordering process will remain unchanged; however, the **review period for swaps** (i.e., the frequency of bid/ask posting) **may occur on a shorter cadence (1-week) to support fast moving deals.**
- **Swaps will be identified and executed through a formalized, central system which will allow for multi-lateral swaps and be managed by Caterpillar.**
  - Dealers will post asks (machines which the dealer wishes to receive in a swap) and bids (current inventory which the dealer is willing to swap away to receive an ask) to a central system on a regular cadence.
  - A central optimization algorithm will identify a series of multi-lateral inventory swaps.
  - These swaps will be communicated to the involved dealers, approved (or vetoed), and executed via Caterpillar's Last Mile network.
  - All differences in the value of machines being swapped will be reconciled according to the wholesale purchase cost of the machines (i.e., if a dealer swaps away a machine worth \$200,000 while receiving a machine worth \$150,000, that dealer will be compensated \$50,000).

## 4. Methodology

The swapping mechanism proposed by this thesis was developed and syndicated in four major phases. First, basic analysis was conducted to prove that virtual pooling would be an impactful tool for the Caterpillar dealer network. Based upon the positive results achieved in that phase, a simplified, demonstration simulation was developed to provide a proof of concept for the network swapping model. This proof of concept was syndicated with a representative subset of dealers to test the concept feasibility and obtain feedback on assumptions and development requirements. Finally, the dealer feedback was incorporated to develop a second-generation, detailed simulation which can be modified to serve as a durable system for identifying swaps.

The research presented in this thesis concluded with the demonstration of a detailed commercial network simulation, and the implications of this simulation are discussed in Results. However, the ultimate objective of this work is a fully-implemented swapping system, and an additional section after the Results is dedicated to the Implementation Pathway.

### 4.1. Preliminary Sales Analysis

Prior to the development of a network swapping mechanism, analyses were conducted to test the hypothesis that virtual pooling (or the approximation thereof) would be an effective tool for reducing held inventory and accelerating sales to end customers. As discussed in Section 3, three conditions must be present in order for virtual pooling to be an efficient tool:

- The **variability in perceived demand at each extant inventory location must be significant**. If demand is highly consistent, the safety stock required to achieve a high customer service level will already be low, and the margin for improvement may not be significant.
- The **demand signals at each inventory location must not exhibit high, positive correlation**. In the event that these demand signals are strongly positively correlated, each point of sale will be attempting to pull on the pooled inventory network at the same time, and little benefit will be realized.
- In the case of virtual pooling, the **logistics cost associated with transshipment must be relatively low**. The benefit of virtual pooling is essentially a one-time inventory reduction (through sales acceleration) less the NPV of the ongoing transshipment logistics costs.

#### 4.1.1. Sales to End User Analysis

The concept of pooling relies on taking variable, uncorrelated demand profiles and leveraging temporally coincident low points one profile and high points on another profile to reduce perceived

variability. Therefore, in order for pooling to be an effective method for managing inventory, the individual demand profiles themselves must have significant variability.

In this case, the relevant demand signal is a customer who is willing to pay for and take receipt of a configured Medium Wheel Loader. Unfortunately, as is typically the case, this data was censored. True demand data has not historically been tracked by most dealers, and it is not tracked with a consistent definition. The most terminal data point which is reliably available is the actual executed sale between the dealer and customer, known as the Sale to End User (STU). One of the assumptions in the research and modeling described in this thesis is that STU data can be used as a proxy for the demand signal perceived by dealers. This assumption was tested during dealer syndication, which is described below, and the feedback collected there suggests that this is a valid assumption.

Using STUs as a proxy for the terminal demand signal, analysis was conducted for the coefficient of variation (CV) in demand at the month / dealer / sales model atom (where  $CV = \text{standard deviation in demand} / \text{mean demand}$ ). In reality, this figure represents a conservative, lower bound on the meaningful CV in demand because it implicitly assumes that each machine within a given sales model is fungible with any other machine in that sales model. This is not the case – as a simple example, standard MWLs designed to move aggregate cannot be substituted for logger machines (which have third-valve functionality, a specialized boom, and an array of additional guarding to protect the machine and operator). The results of this analysis are summarized in Figure 6 below.

### Histogram of Observed Coefficient of Variation in Demand

n = number of month / dealer / sales model atoms in CV range

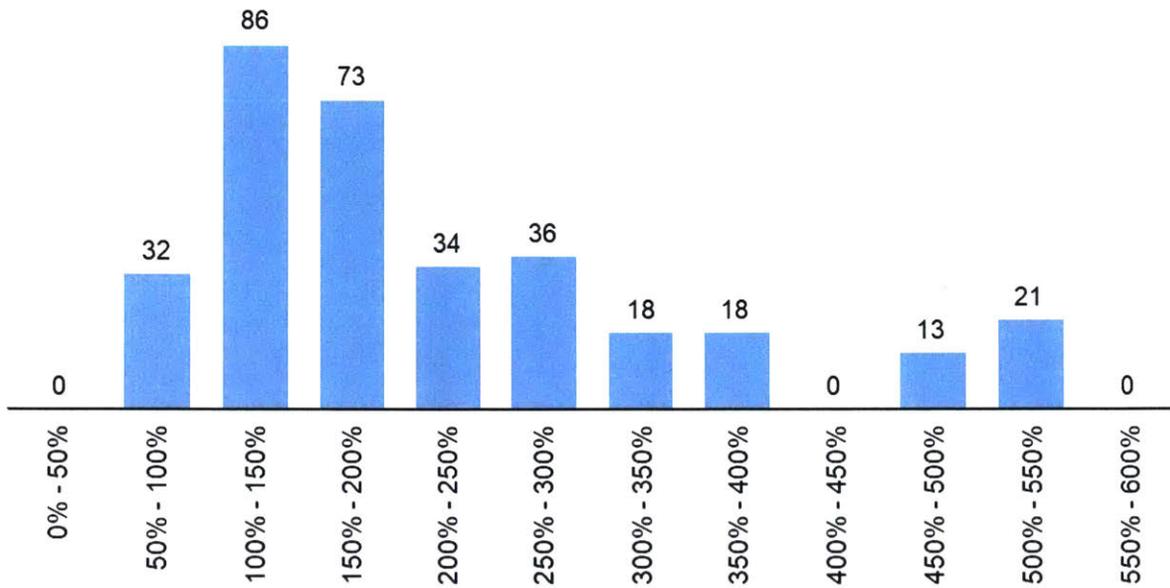


Figure 6: Histogram of observed CV at a month / dealer / sales model level

The minimum observed CV in demand was 55%, while the maximum observed CV exceeded 500%; the unit-weighted average observed CV in demand at the month / dealer / sales model level was 136%, and the median CV was 181%. Figure 7 presents a disguised monthly STU series for a single sales model at a single dealer; the CV of the series presented is 107%. With the exception of Caterpillar's largest dealers in ADSD-N, this pattern – where a dealer will make several sales in a month followed by a series of months with no sales – is typical.

## Example Monthly Sales Series

Sales to End User (Example dealer / sales model cross section)

---

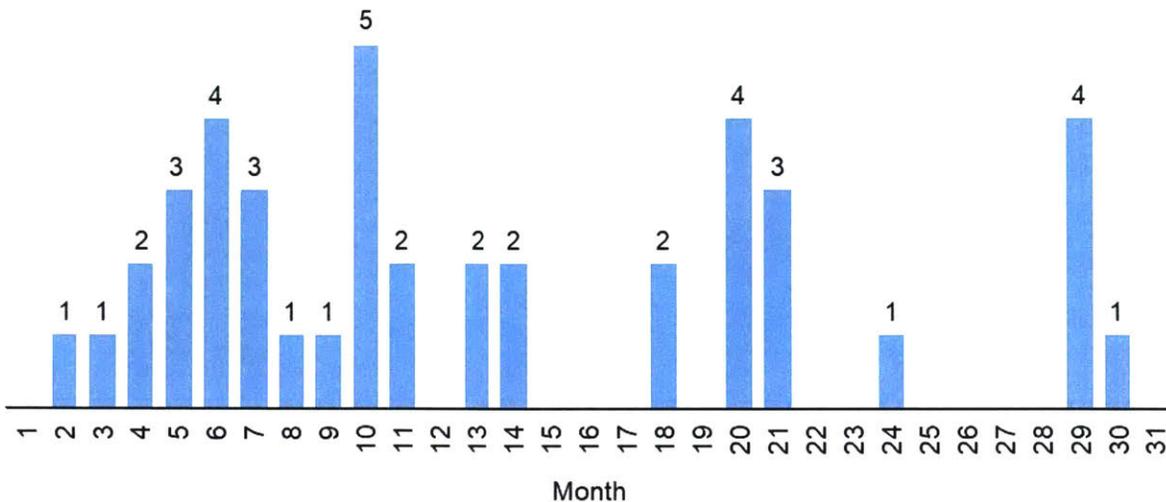


Figure 7: Example Sales to End User series

Another important prima facie implication of this demand analysis is that normality is not a valid assumption. Monthly demand is low enough that it should be treated as a discrete distribution; furthermore, each demand distribution will display significant positive skewness (due to the frequency of months with no demand).

A simple analytical model can be created to impute the approximate  $\beta$  service level implied by each dealer's historical stocking level ( $\beta$  service level is defined here as the proportion of *units* across the demand series which are served from inventory; by contrast  $\alpha$  service level is the proportion of periods in which all orders can be served from inventory). This simplified model makes the following assumptions:

- Factory lead time is fixed at three months (i.e., all machines are ordered through Lane 3, and the Caterpillar factory has sufficient capacity such that lead time does not flex)
- Dealers follow a simple base stock re-ordering algorithm with a review period of one month (i.e., dealers have no insight into forthcoming deals; forward coverage reordering is not used)
- Customers are willing to wait indefinitely to receive their machine; in the event of a stock-out they will wait to receive that machine until it is in stock
- Backorders are served prior to new orders

In the example sales series above, the dealer maintained a historical average inventory of approximately 1.83 machines, and the mean sales per month were 1.35 machines. If we make the additional

assumption that dealer stock is composed exclusively of cycle stock and safety stock, we can calculate the dealer's target base stock by first calculating the safety stock and then using this figure to calculate the base stock (Note:  $r$  = review period,  $L$  = factory lead time,  $\mu$  = mean monthly demand):

$$\text{Avg. inventory level} = \text{cycle stock} + \text{safety stock}$$

$$\text{Avg. inventory level} = \frac{r\mu}{2} + \text{safety stock}$$

$$1.83 = \frac{1 \cdot 1.35}{2} + \text{safety stock}$$

$$\text{safety stock} = 1.16 \text{ machines}$$

$$\text{Base stock} = (r + L)\mu + \text{safety stock}$$

$$\text{Base stock} = (1 + 3)1.35 + 1.16$$

$$\text{Base stock} = 6.58 \text{ machines}$$

Based upon the above calculated, an 'order to' figure of seven machines was set, and the simple model was constructed to observe the expected resulting service level. Figure 8 shows the results of the simple model for the example case presented in Figure 7 (model details are presented in Appendix: Service Level Sample Analysis).

### Example Monthly Order Service Series

Sales to End User (Example dealer / sales model cross section)



Figure 8: Example simple service level model results

In this example, the  $\beta$  service level is 83% (35 orders served immediately out of 42 total orders). Across all dealers and models, the average unit-weighted modeled  $\beta$ -service level was 90.6%. In order to

maintain this service level, however, Caterpillar's dealers must maintain a prodigious level of safety stock – these calculations suggest that, on average, the dealer network holds more than seven units of safety stock for each unit of cycle stock.

Contextualizing the implications of these stocking levels requires an understanding of how Caterpillar and its dealers view inventory holding costs. From a financial perspective, Caterpillar treats inventory the same way as its other fixed and variable assets, applying the same capital charge (which is intended to represent the company's true cost of capital) that it uses for OPACC calculations. Dealers generally assume a much lower financial holding cost (approximately equivalent to low-risk business loan rates) but also must account for the direct costs of exercising and maintaining equipment on their lots.

The above analysis ignores certain real-world characteristics of Caterpillar's commercial network. For example, dealers maintain close relationships with major customers in their geographic territory, and they likely have some capacity to forecast when large orders will arrive. This foresight would allow the dealer to build anticipatory stock for significant order events, possibly improving their real-world service level. However, both the CV analysis and simple service level model corroborate anecdotal evidence provided by Caterpillar's dealers: demand variability is high, and dealers must either sacrifice service level or carry high levels of inventory to meet potential customer demand.

#### 4.1.2. Demand Correlation

The presence of high demand variability and high safety stock (or depressed service levels) is necessary for pooling to be effective, but it is not sufficient. The variable demand signals at each point of sale must also be relatively uncorrelated. The reason for this is simple: if demand for a given sales model (or all sales models) peaks across the entire network at the same time, then there is no surplus stock to be shared.

Interviews conducted within Caterpillar and at its dealers initially suggested that the level of demand correlation across dealers may be high. The dealers interviewed suggested that when their sales peaked, so did the sales of neighboring dealers; informal inventory trading which may occur during 'typical' demand periods would slow down as dealers sought to hold on to every unit of inventory which represented a potential sale. Individuals within Caterpillar itself also noted the presence of 'hoarding' behavior that arises when overall demand rises across its entire system of dealers, and it has put measures in place to counter this behavior.

Although demand correlation likely does appear when examining longer time horizons (e.g., across economic cycles), empirical evidence suggests that demand correlation among dealers at the monthly time scale relevant for finished goods inventory pooling is low. Again using Sales to End Used data as a

proxy, the demand correlation for each pair of dealers in the ADSD-N territory was examined across all eight Medium Wheel Loader sales models. Correlation here is defined as the Pearson product-moment correlation, given by the equation below.

$$r = \frac{\sum(x - \bar{x})(y - \bar{y})}{\sqrt{\sum(x - \bar{x})^2(y - \bar{y})^2}}$$

The results of the demand correlation analysis are summarized in Figure 9.

Sales Model	Intra-District Demand Correlation	Extra-District Demand Correlation	Overall Demand Correlation
950M	0.089	0.048	0.057
962M	(0.002)	0.033	0.026
966M	0.055	0.036	0.041
966MXE	0.007	0.001	0.002
972M	0.001	0.018	0.014
972MXE	0.047	0.002	0.012
980M	0.070	0.070	0.070
982M	0.076	0.062	0.065
<b>Overall MWL</b>	<b>0.062</b>	<b>0.047</b>	<b>0.051</b>

Figure 9: Dealer to dealer demand correlation summary

In the table above, Intra-District demand correlation refers to the demand correlation between dealers that are located in the same marketing district. Each of the eight districts is composed of approximately five to eight dealers which are geographically adjacent to each other and interact with end users in a similar distribution of industries (see Appendix: Dealer District Map to see a map of ADSD-N dealer districts). The relevance of dividing this analysis between intra-district and extra-district correlations is detailed in the Implementation Pathway segment: the proposed pilot swapping pools will occur within individual districts.

Overall, demand correlations observed at the intra-district and extra-district level are very slightly positive. This suggests that there may be some underlying factors which drive a portion of demand across many dealers (e.g., customers may wish the purchase vehicles prior to the end of a fiscal year), but these effects are relatively small.

Examining demand variance at the dealer and district level provides an alternative avenue to estimate the impact of demand correlation on pooling efficacy. If we make the assumption that each dealer’s perceived demand can be represented as a bundle of independent Poisson distributions, then the overall demand observed at a dealer or district level is linearly related to the number of bundled Poisson distributions observed by that dealer or district. Consequently, we would expect a plot of the coefficient

of variation in demand versus average monthly demand to yield a power function with an exponent of -0.5. Figure 10 shows a plot of the empirical data for a single sales model with the fitted power function.

**Coefficient of Variation vs. Mean Monthly Demand**

% CV vs. Units per month (Example sales model)

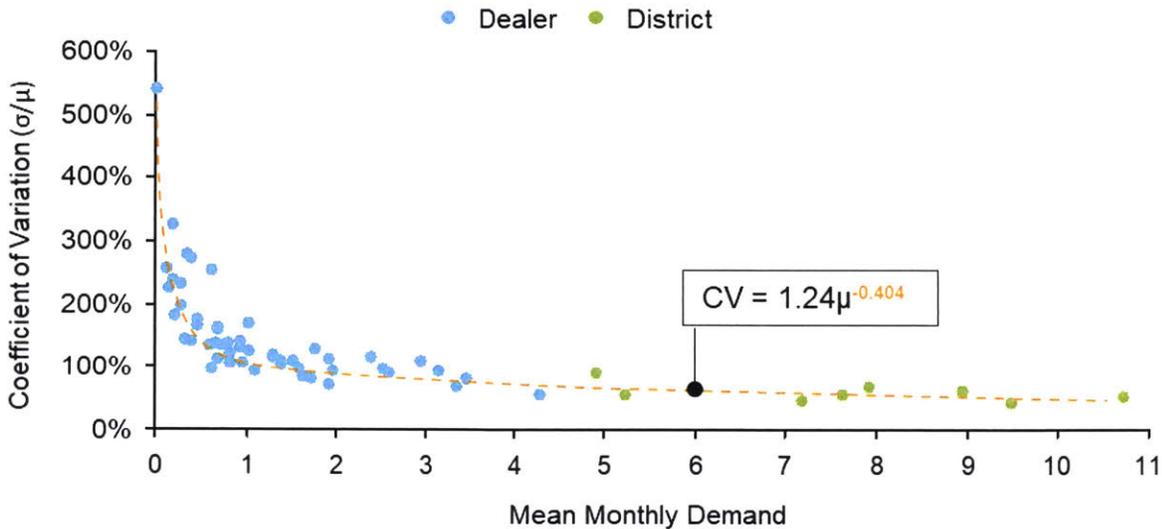


Figure 10: Coefficient of variation vs. mean monthly sales, example sales model

Across all models, the observed fitted  $\lambda$  ranged from -0.404 (for the sales model shown in Figure 7) to -0.483. The fact that this coefficient is slightly lower in magnitude than the result proposed by Eppen in 1979 can possibly be explained by the fact that the asymmetric demand distribution has somewhat heavier tails than a normal distribution. [22] However, taken in conjunction with the observation that demand correlation across dealers is low and the observed coefficient of variation across dealers is high, this coefficient is sufficiently high to suggest that inventory pooling is likely to be an effective mechanism for reducing safety stock or improving service level.

4.1.3. Virtual Pooling Opportunity

Using Sales to End User Data, a simple model was developed to examine the theoretical impact of virtual pooling without the dealer ownership constraint addressed in this thesis.

The model extends the simple base-stock model used to estimate  $\beta$  service level (described in Section 5.1.1), and the same underlying assumptions apply. In addition to the individual dealer-by-dealer base stock calculations, the model applies the same calculations to a demand series equivalent to the sum of the observed demand series at each dealer within the marketing district. The assumption implied by this calculation is that dealers within a given marketing district freely transfer inventory to each other in the event that any dealer in the network would otherwise face a back order and other dealers have

inventory on hand which can meet the appropriate order requirement. Sales model variants are not considered. In reality – even if the ownership agency problem is ignored – a true virtual pool would be more constrained than this model suggests, but it does provide a first approximation of the potential pooling benefit.

At each target inventory position for which the  $\beta$  service level was calculated, an optimization was conducted to determine which dealer in the marketing district should hold that inventory in order to maximize the overall service level across dealers. The formulation for this optimization is presented below.

*The monthly commercial cycle for a single dealer in the specific case of three-month lead times can be represented by the following system of equations:*

$a_t = i_{t-3}; a_0 = 0, a_1 = 0, a_2 = 0$	Factory deliveries to dealer in period t
$b_t = h_{t-1} + a_t; b_0 = x$	Period beginning inventory
$d_t = \min(b_t, g_{t-1})$	Sales against backorders
$e_t = \min(b_t - d_t, c_t)$	Sales against current orders
$f_t = \max(0, c_t + d_t - b_t)$	New backorders
$g_t = f_t - d_t + g_{t-1}; g_0 = 0$	Period ending backorders
$h_t = b_t - d_t - e_t$	Period ending inventory
$i_t = \max(0, x - h_t - i_{t-1} - i_{t-2} + g_t)$	Dealer restock orders to factory

*Where  $c_t$  is the known customer demand in each period,  $x$  is a variable defining the target inventory position for the dealer, and  $t$  is an integer ranging from 1 to 48 representing the month of the simulation.*

*This system of equations can be reduced to the following:*

$$a_t = \max[0, x - b_{t-3} + \min(b_{t-3}, g_{t-4}) + e_{t-3} - i_{t-4} - i_{t-5} + g_{t-3}]; a_0 = 0, a_1 = 0, a_2 = 0$$

$$b_t = b_{t-1} - \min(b_{t-1}, g_{t-2}) - e_{t-1} + a_t; b_0 = x$$

$$e_t = \min(b_t - \min(b_t, g_{t-1}), c_t)$$

$$g_t = \max(0, c_t + \min(b_t, g_{t-1}) - b_t) - \min(b_t, g_{t-1}) + g_{t-1}; g_0 = 0$$

*Once this system of equations has been established, the optimization to distribute inventory across a set of dealers to maximize  $\beta$  service level may be written as:*

$$\max \sum_{d \in D, t \in T} e_{dt}$$

*Subject to:*

$$\sum_{d \in D} x_d \leq X$$

$$x_d \geq 0$$

Where  $e_{dt}$  are the sales made against new customer orders for a single period at a single dealer belonging to the ordinal set of dealers  $D$ . The overall target inventory position,  $X$ , was progressively incremented by one unit from zero until both the pooled and unpooled case demonstrated a 100% service level, and the optimization described above was conducted at each target inventory position. The  $\beta$  service level was observed at each point and plotted against the average on-hand inventory. An example of this plot for a single sales model and marketing district is presented in Figure 11.

### **$\beta$ -Service Level vs. Average On-Hand Inventory, Example Model and District**

% of orders served immediately vs. average units of inventory

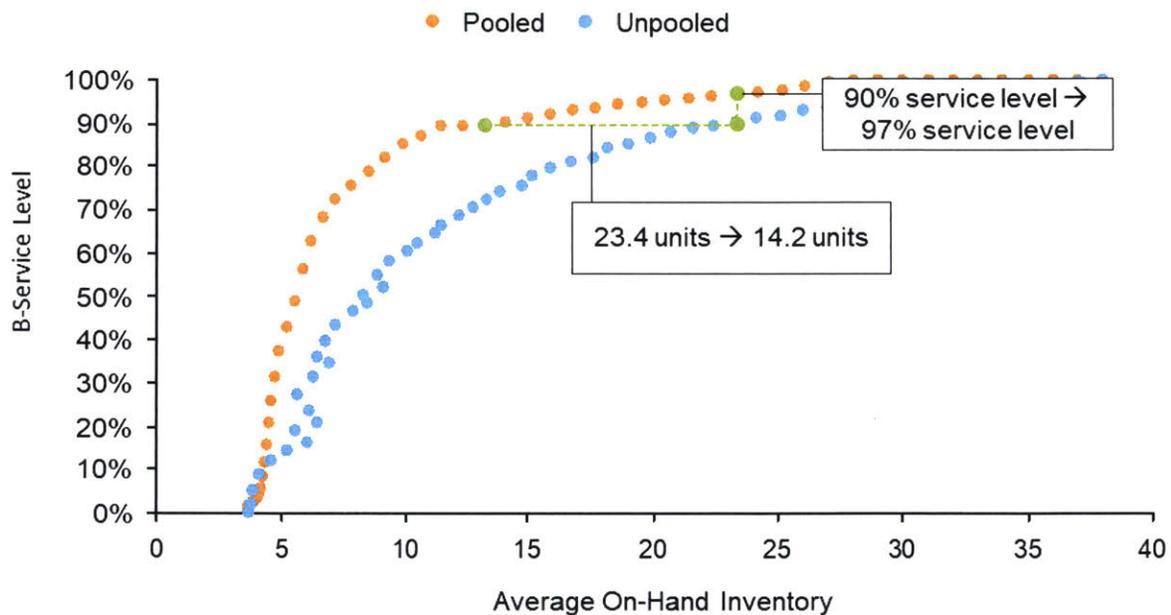


Figure 11: Preliminary virtual pooling model results, example model and district

In the example above, the marketing district requires an average on-hand stocking level of 23.4 units in order to achieve a 90% service level (the average ADSD-N, Medium Wheel Loader service level calculated in Section 5.1.1) without inventory pooling. Under the pooled condition, the district is able to achieve a service level of 97% at a stocking level of 23.4 units. Conversely, it could reduce average on-hand inventory to 14.2 units while maintaining a 90% service level.

It is important to note that this analysis is not financial and does not address logistics cost, which is a major driver of the financial impact of virtual pooling. In the case of Caterpillar’s general Construction Industries group, however, transportation costs are small relative to the cost of the equipment (on the

order of 1 to 5% of the equipment's sale price), and the margins realized by dealers upon making a sale are relatively large. Qualitatively, at least, this suggests that the increased logistics cost imposed by virtual pooling will be outweighed by the benefit of accelerated sales and lower inventory holding costs.

The proposed swapping mechanism developed in section 5.2 is significantly more constrained than the virtual pool modeled above (for example, machine configurations must be matched to some degree in real-world swapping), and the commensurate benefit will be lower than this analysis suggests. However, these results are compelling enough to justify the exploration of a swapping network which could capture some of the benefits of a true virtual pool.

#### 4.2. Network Simulation Development

The swapping network simulation discussed in this thesis was developed in two generations. A preliminary, limited simulation was developed in Excel and Visual Basic to develop the basic model structure and assumptions. The results from this model were syndicated with dealers (discussed in section 5.3 below), and the feedback was used to develop a durable model in Python which was used to develop the Results discussed in this thesis.

##### 4.2.1. Overall Network Swapping Model Structure

The network swapping model developed in this thesis is intended to simultaneously simulate the Caterpillar commercial network under conditions with and without network swapping in order to observe the potential improvement in inventory management and commercial velocity.

The overall structure of the model is presented in Figure 12 below.

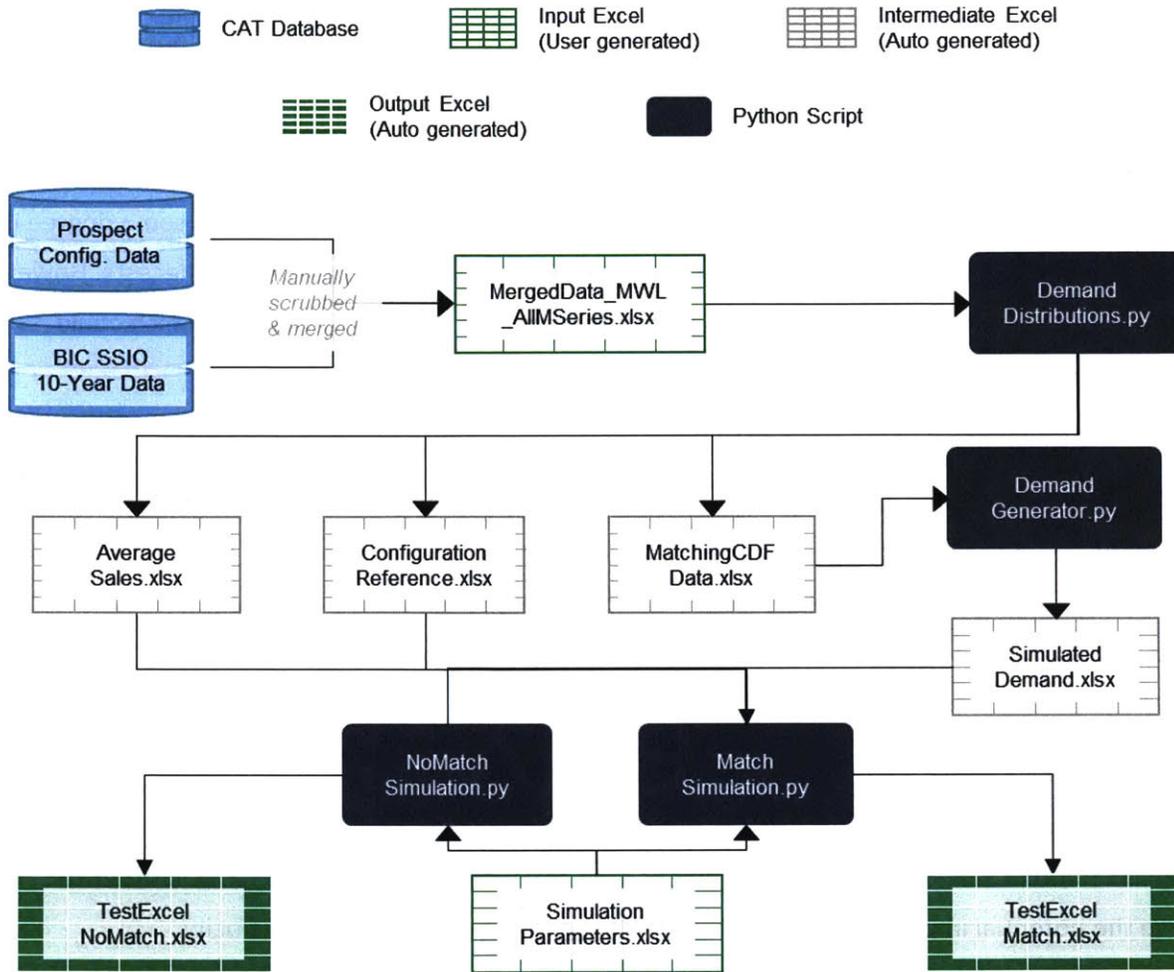


Figure 12: Overall simulation model architecture

The simulation is designed in several modules so that multiple parameter sets can be run against the same modeled demand series and to facilitate a Monte Carlo simulation.

The model is designed to simulate the Caterpillar commercial network for up to a 48-month long period with a monthly time step. Figure 13 below presents the commercial processes that the simulation runs through during each period.

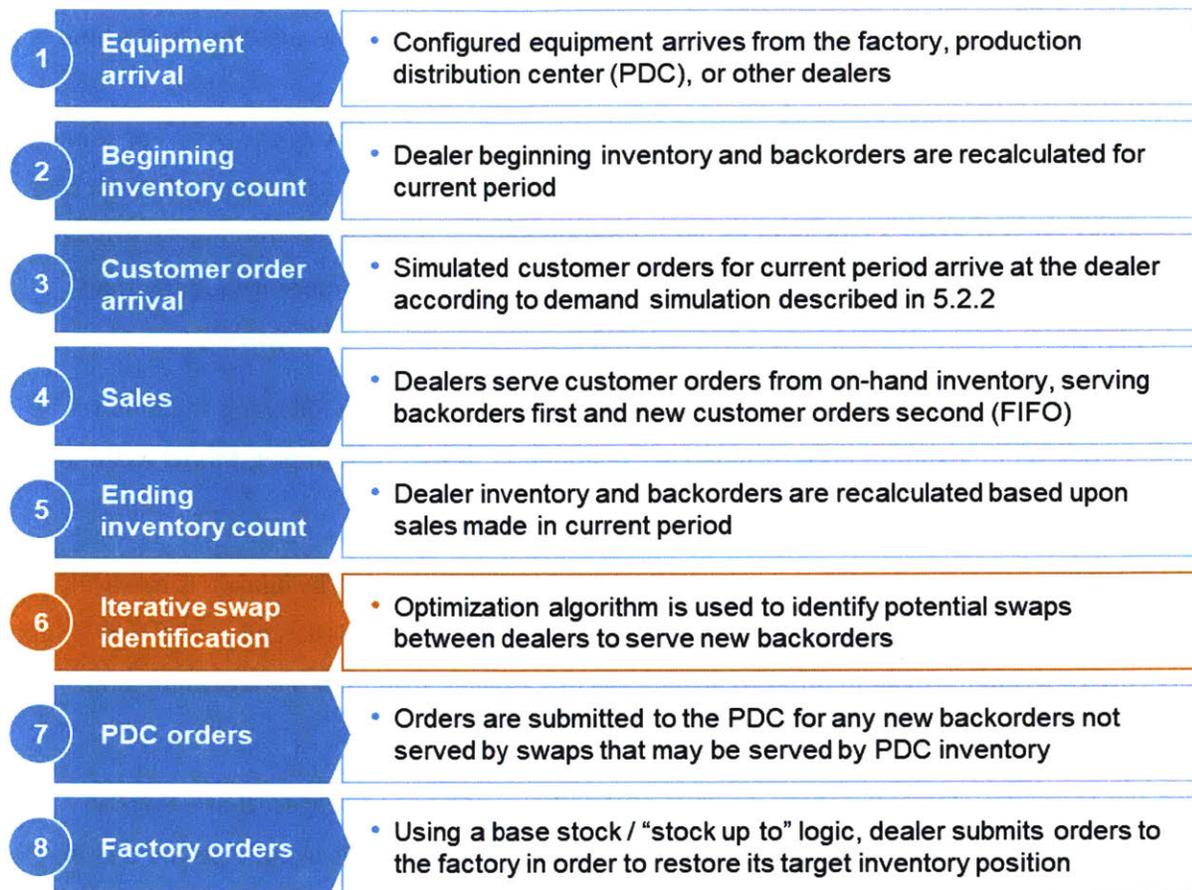


Figure 13: Simulated period commercial process order

The sixth process – iterative swap identification – is idiosyncratic to the Match Simulation; otherwise, the Match and No Match simulation follow the same commercial process order. The commercial process outlined above is a simplified model of the actual processes faced by Caterpillar’s dealers, and there are a number of assumptions implicitly made in this modeled process:

- The model assumes that there is **distinct point in time at which a customer order “arrives” at the dealer**; the order arrival is the point at which a customer is willing to pay for and take receipt of a configured machine. While this discrete order arrival does occur in many cases, some orders are the culmination of ongoing conversations where no clear order event arises – in this latter case, customers may be less sensitive to the timing of machine delivery, potentially obviating the need for swaps.
- In the case that it faces a backorder, the **dealer tries to source machines according to the following primacy: (1) it solicits a swap (2) it orders a machine from the PDC and (3) barring the other two options, it orders a new machine from the factory**. The justification for this assumption is that the dealer would wish to serve the backorder as quickly as possible; however,

the dealer may elect to serve a backorder from the factory first if the customer is not time sensitive.

- The model assumes that **all of the above processes occur at a monthly cadence**. For some processes, such as factory orders and deliveries, this is an accurate assumption. Others, such as order arrivals, occur on an ongoing basis. Consequently, subsets of the commercial process described above (e.g. order arrivals and swap sourcing) may occur more frequently than monthly.

Additional critical assumptions made in this simulation are discussed in the following sections. A comprehensive set of assumptions is included in the Appendix: Network Swapping Model Assumptions.

#### 4.2.2. End User Demand Modeling

As discussed in section 5.1.1, the Sales to End User (STU) demand observed by dealers is relatively low and exhibits significant positive skewness. Furthermore, the network simulation discussed in this section requires discrete order arrivals in order to analytically model sales, inventory management, and swapping.

As discussed in section 5.1.2, the characteristics of the observed demand suggest that a Poisson distribution may be an effective descriptive model. In order to examine this hypothesis, a Poisson distribution was fitted to each demand series at the dealer / sales model level, and a chi-squared goodness of fit test was conducted on the resulting fitted distributions according to the following equations. [23]

The parameter lambda was estimated according to:

$$\lambda = \frac{\sum_{j=1}^c d_j f_j}{n}$$

Where c is the maximum number of units demanded in one month, d is an integer corresponding to a given level of demand, f is the frequency with which that level of demand was observed, and n was the total number of months in the demand series.

The chi-squared statistic was calculated according to:

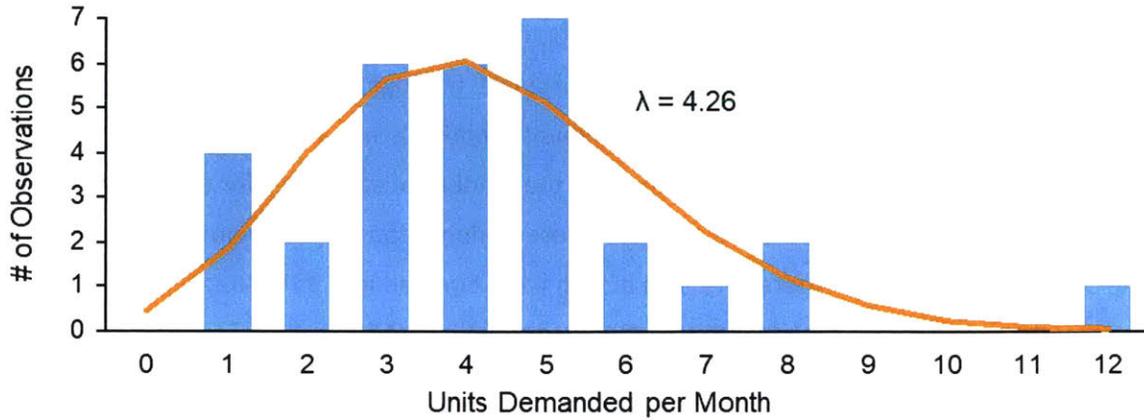
$$X_{k-2}^2 = \sum_k \frac{(f_o - f_e)^2}{f_e}$$

Where k is a given level of monthly demand,  $f_o$  is the observed frequency at a given level of monthly demand,  $f_e$  is the predicted frequency at a given level of monthly demand, and  $k - 2$  is the number of degrees of freedom of the test (the number of parameters estimated from the data is one).

The very low demand observed across some dealers and models as well as the limited number of observations disqualified many of the demand series from a chi-squared test. Across the 106 series for which a chi-squared goodness of fit test could reasonably be conducted, 34 were significant at the  $p = 0.05$  level (i.e., the null hypothesis was rejected). It is unlikely that the monthly demand variance observed in these cases could be fully explained with a single, homogenous Poisson distribution. Histograms of observed demand series for which the null hypothesis was and was not rejected are presented in Figure 14 below. It should be noted that the number of categories for the chi-squared test was circumscribed at the point where the predicted observation frequency fell below one. In the top case, the observations with 7 or more sales in a month were bundled into a single category; in the bottom case, the observations with 6 or more sales were bundled into a single category.

**Demand Histogram, Example Model / Dealer,  $p = 0.500$  (Null hypothesis NOT rejected)**

Observation frequency vs. monthly units demanded



**Demand Histogram, Example Model / Dealer,  $p = 0.010$  (Null hypothesis rejected)**

Observation frequency vs. monthly units demanded

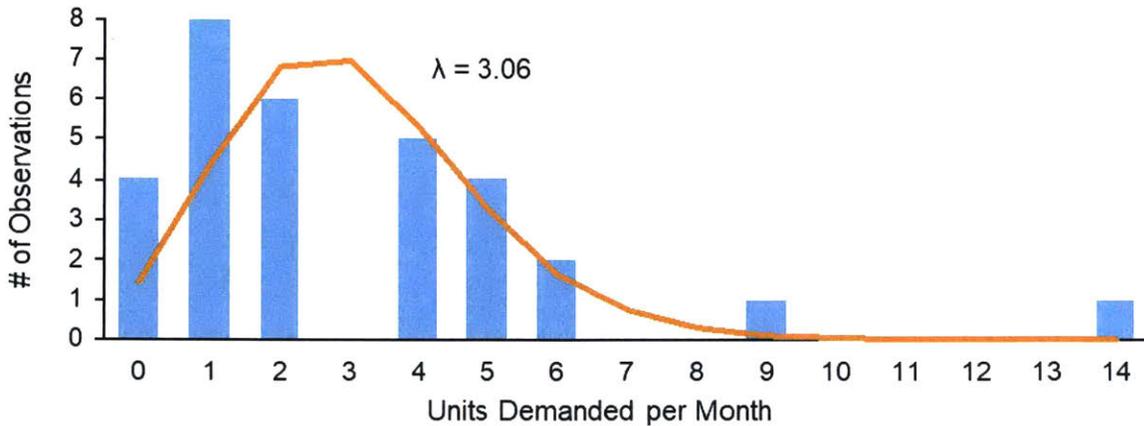


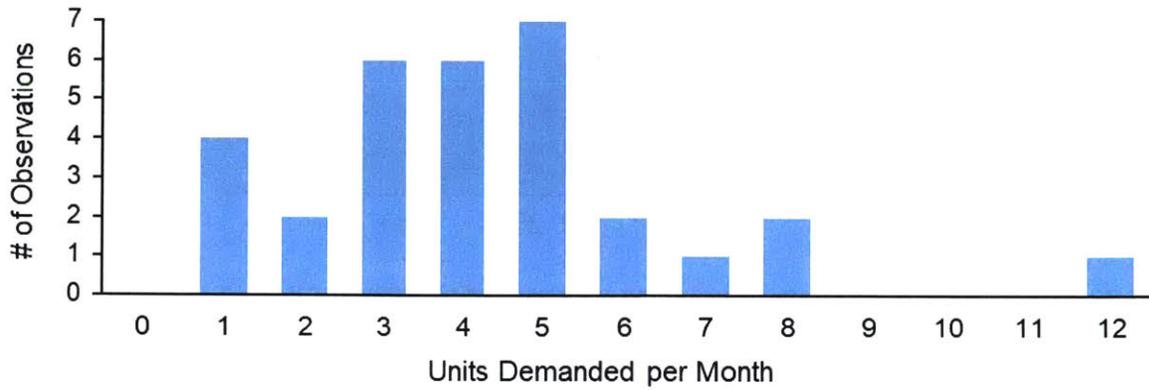
Figure 14: Demand histograms with fitted Poisson distributions

Given the observation that a Poisson distribution is unlikely to be an effective stochastic model in at least some of the dealers and sales models, the decision was made to use historical STU data to construct an empirical cumulative distribution function (CDF) for each dealer and sales model. When creating a synthetic demand series, this CDF was sampled using the uniform distribution  $U(0,1)$  for each month of the modeled series. The process of generating one such empirical CDF and sampling the resulting CDF are shown in Figures 15 and 16 below.

### Demand Histogram

Observation frequency vs. monthly units demanded

---



### Cumulative Distribution Function

Cumulative probability vs. monthly units demanded

---

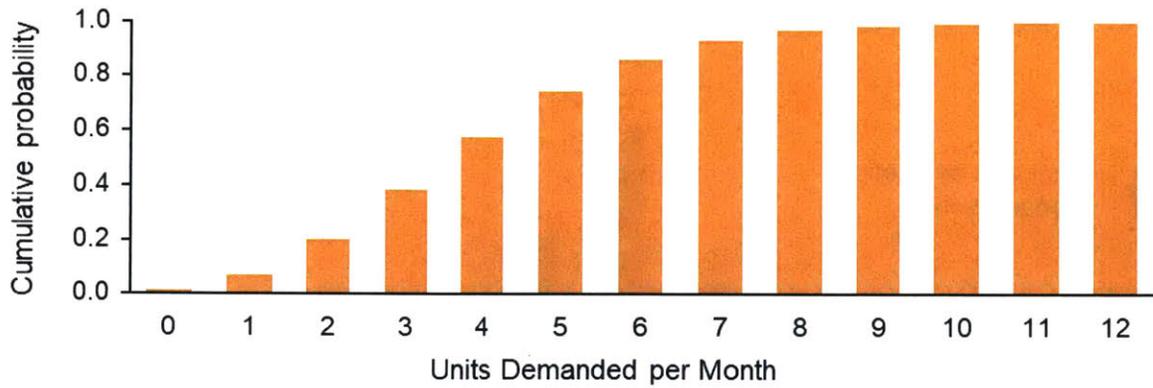


Figure 15: Example empirical demand cumulative distribution function

## Empirical CDF Sampling

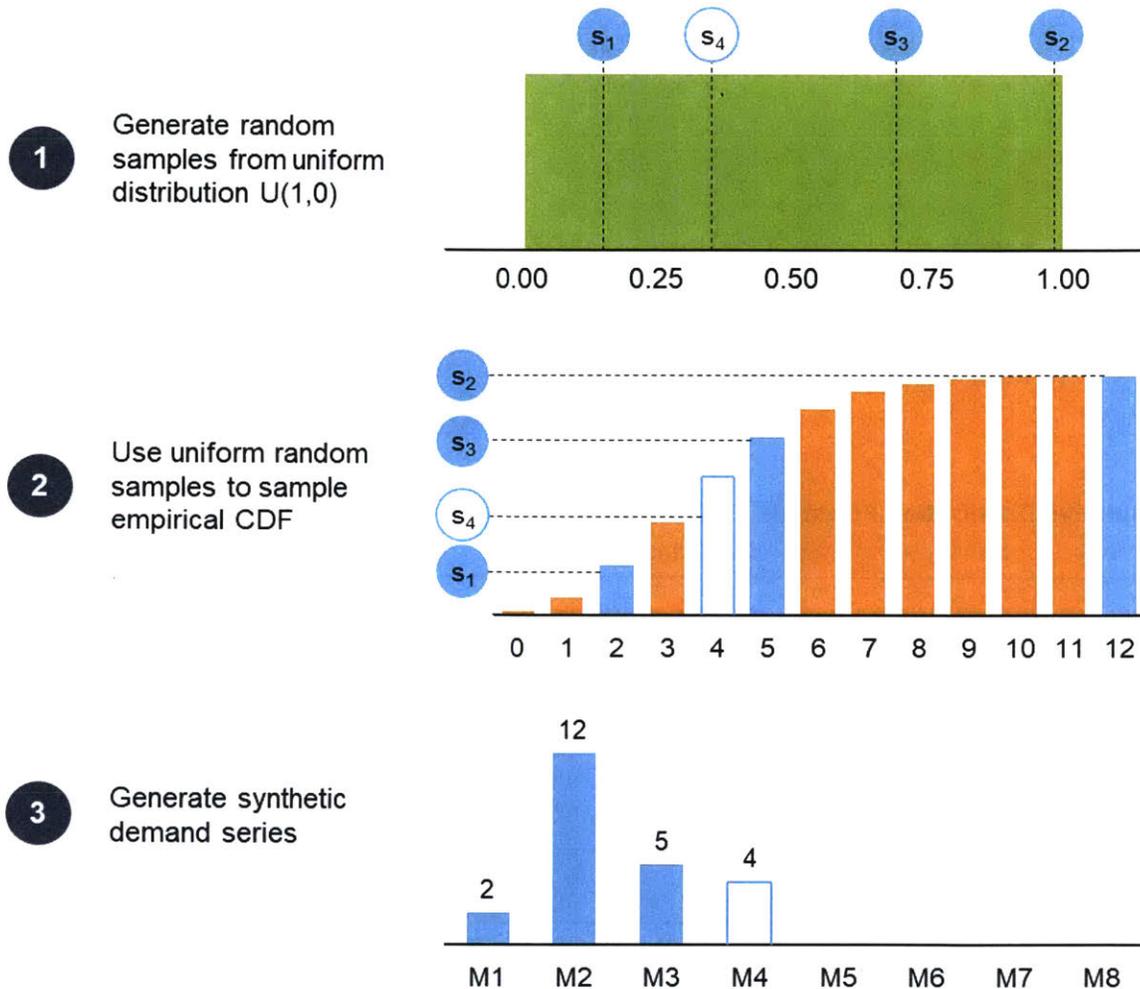


Figure 16: Generating a synthetic demand series from the empirical CDF

While the first generation model did not account for configuration complexity, the second generation model uses a similar method to assign a desired configuration to each simulated order. The configurations of historical orders were examined, the frequency of each configuration was determined, and a separate cumulative distribution function was generated for each dealer / sales model atom. During the generation of a synthetic demand series, this CDF is sampled independently from the one which determines the overall level of demand (Note: this implicitly assumes that the configurations of machines that are ordered by a dealer in a given month are not correlated. This is a simplifying assumption for the sake of simulation and is likely violated in reality if dealers are anticipating sales to customers that demand a certain type of configuration). It's important to note that this methodology is limited by the fact that the data history for the current generation of wheel loaders was limited to approximately three years, and long term patterns in demand may not be sufficiently stabilized.

Depending upon which parameters are defined by the user as 'significant,' these full complexity configurations are merged into the subset of potential configurations that can be represented by the subset of configuration dimensions. Figure 17 below provides an illustrative example of this configuration reduction. This compression is only conducted during the simulation and does not affect the simulated demand. Consequently, information about configuration is not lost, and multiple simulations can be run against the same modeled demand profile.

**Example simulated demand, full configurations**

Boom	Hydraulics	Fenders	Differential	Lights	Config #
Standard	3-valve	Standard	Manual	Halogen	1
Standard	2-valve	Standard	Manual	Halogen	2
Standard	3-valve	Roading	Manual	LED	3
Standard	2-valve	Standard	Auto	LED	4
Standard	2-valve	Roading	Manual	Halogen	5

**Example simulated demand, simplified configurations**

■ 'Significant' configuration parameters

Boom	Hydraulics	Fenders	Differential	Lights	Simple Config #
Standard	3-valve	<i>Standard</i>	Manual	<i>Halogen</i>	1
Standard	2-valve	<i>Standard</i>	Manual	<i>Halogen</i>	2
Standard	3-valve	<i>Roading</i>	Manual	<i>LED</i>	1
Standard	2-valve	<i>Standard</i>	Auto	<i>LED</i>	3
Standard	2-valve	<i>Roading</i>	Manual	<i>Halogen</i>	2

Figure 17: Example configuration compression; five unique configurations are reduced to three by identifying three 'significant' configuration parameters

### 4.2.3. Inventory Management Modeling

The first generation model included two inventory modeling methodologies: forward-coverage and base stock. In practice, most dealers use a system called DDVS, which is based upon forward-coverage but allows for significant subjective adjustments by the dealers themselves.

The preliminary modeling suggested that the selection of a forward-coverage or base stock methodology had little impact on the relative benefit associated with network swapping. Furthermore, Neale and Williams have demonstrated that forward-coverage models can result in a phenomenon known as the ‘landslide effect’ in the presence of seasonal demand and underperform base stock models. [24] Consequently, the second generation model assumes a simple base stock, ‘order up to’ methodology of inventory control. Under the no-swapping condition, dealers preferentially order from the Product Distribution Center (PDC) and then fill their remaining orders from the Caterpillar factory. Consequently, the no-swap factory order for period  $t$  can be written as:

$$R_{f,t} = C_t - R_{pdc,t}$$

Where  $R_{f,t}$  is the order submitted to the Caterpillar factory during period  $t$ ,  $C_t$  is the total customer order arrivals in period  $t$ , and  $R_{pdc,t}$  is the order submitted to the PDC in period  $t$ .

In the swapping case, the reorder logic also accounts for swaps that the dealer has committed to. Consequently, the swap condition factory order for period  $t$  can be written as:

$$R_{f,t} = C_t - R_{pdc,t} - S_{+,t} + S_{-,t}$$

Where  $S_{+,t}$  are the machines that the dealer will be receiving via committed swaps and  $S_{-,t}$  are the machines that the dealer will be sending to other dealers via committed swaps.

In some cases, product at the PDC will be constrained (e.g., the PDC may only hold five units of a specific model and configuration, but the period customer orders across all dealers will exceed five). In this case, the PDC’s inventory is sequentially allocated to the requesting dealers one at a time until the PDC inventory has been exhausted. In reality, PDC inventory is typically allocated based upon a combination of perceived dealer need and a first-come, first-served methodology; however, this sequential allocation helps ensure that the model does not arbitrarily generate results that are preferential to a dealer or set of dealers.

The review period in the model is fixed at one month, while the factory lead time can be adjusted to model different levels of system constraint. The network swapping model makes the simplifying assumptions that factory lead time remains constant (i.e., the factory constraint is not a function of

intra-simulation order quantities) and does not vary between Lane 2 and Lane 3 orders (which will have different lead times under most conditions).

#### 4.2.4. Swapping Identification

In order to identify feasible swaps, the swapping simulation sets up and solves a series of network optimization problems during each period of the simulation. The algorithm itself is a simple, unweighted, undirected longest path problem which may or may not be acyclic depending upon the presence of an altruistic donor (the Product Distribution Center, or PDC).

An illustrative example is helpful in describing the formulation of the optimization problem that the swapping model solves. Consider the following situation with a three dealer network which is selling three sales models, each of which has a single possible configuration. The dealers have concluded the sales cycle of period  $t$ , and each dealer has two items of remaining inventory and two backorders that it was unable to serve from its own inventory. This situation is presented in Figure 18.

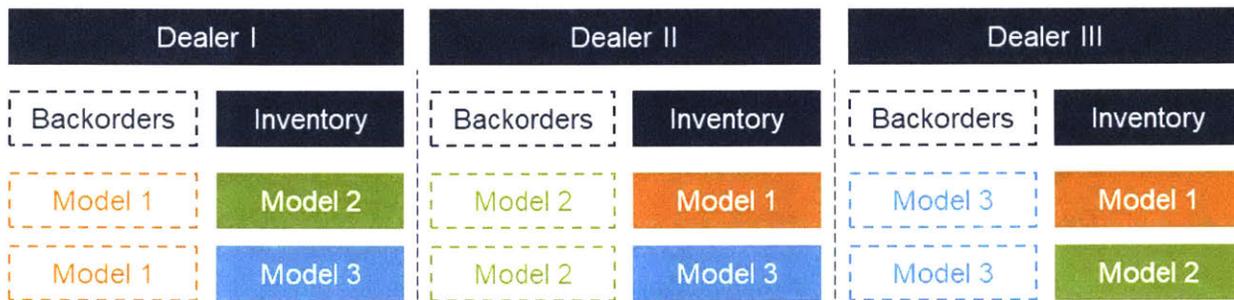


Figure 18: Illustrative swapping optimization scenario

The first step in formulating the associated optimization problem is to draw the network across which the optimization will occur. In this case, there are three nodes (the dealers) with six possible directed edges between each pair of nodes (one edge for the possibility of sending or receiving each machine). Since a given pair of nodes may be connected by multiple directed edges, this network is considered a quiver. The network for the illustrative example is presented in Figure 19.

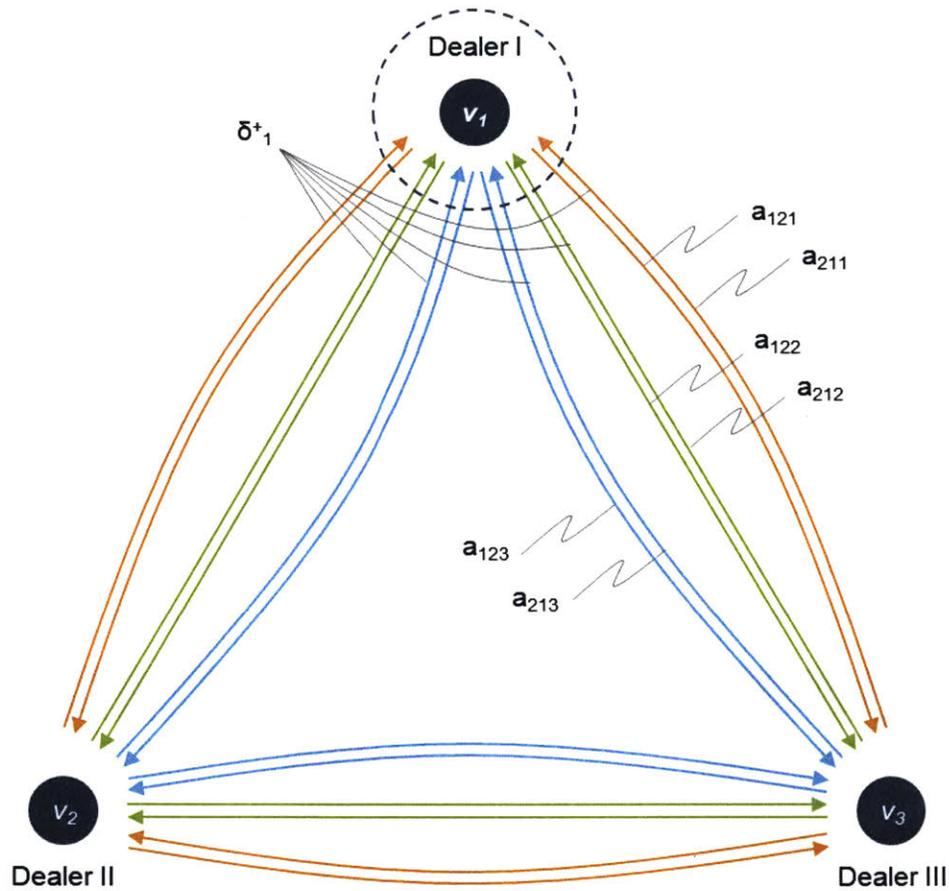


Figure 19: Illustrative swapping network

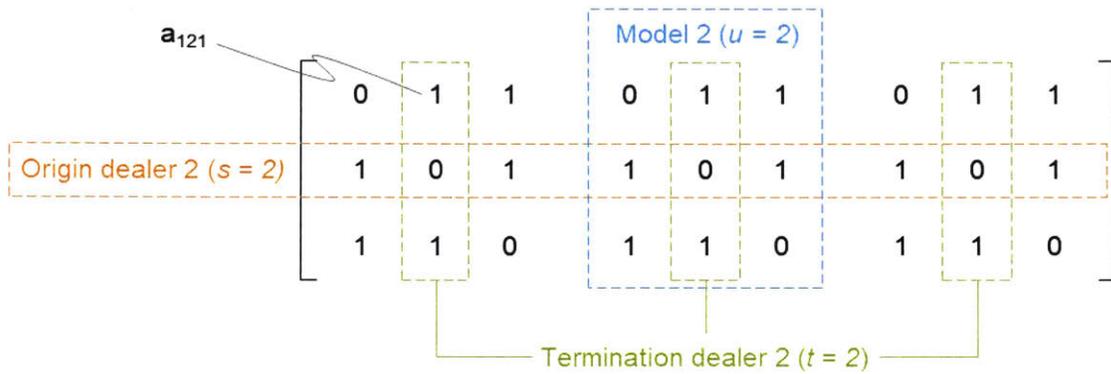
The quiver can be denoted as  $G(V,A,s,t)$  where  $V$  is the set of vertices (dealers) indexed by  $i$ , and  $A$  is the set of arcs indexed according to three integers:  $s$  (the origin vertex/dealer),  $t$  (the termination vertex/dealer), and  $u$  (a unique model/configuration). For example, the arc  $a_{121}$  in the network above is the edge directed from Dealer I to Dealer II along which Machine 1 can be sent. In order to simplify constraint formulation, the full set of incoming arcs to a given vertex,  $i$ , is called  $\delta_i^+$ , and the full set of outgoing edges from a given vertex,  $i$ , is called  $\delta_i^-$ . [25]

Before conducting the optimization, the three-dimensional matrix  $A$  representing the arcs in the network is adjusted by taking the elementwise product with two additional matrices: a durable matrix  $B$ , which defines any ongoing transportation constraints, and a temporary matrix  $C$ , which adjusts graph arcs based upon the coincident presence of a unit of inventory at the origin of an arc and the presence of a backorder at the termination of an arc.

In the illustrative example above, some ongoing transportation constraint may prevent Dealer I from shipping machines to Dealer II. A real world example of what could cause such a transportation constraint is the presence of an international border which makes swapping infeasible or a geographic

distance which makes swapping cost prohibitive (e.g., dealers in Maine and California may not wish to swap directly with each other). In the case of Dealer I not being able to ship machines to Dealer II, the matrix A is adjusted by taking its elementwise product with matrix B, as shown in Figure 20 below.

**Matrix A (full network graph)**



**Matrix B (durable transportation constraints)**

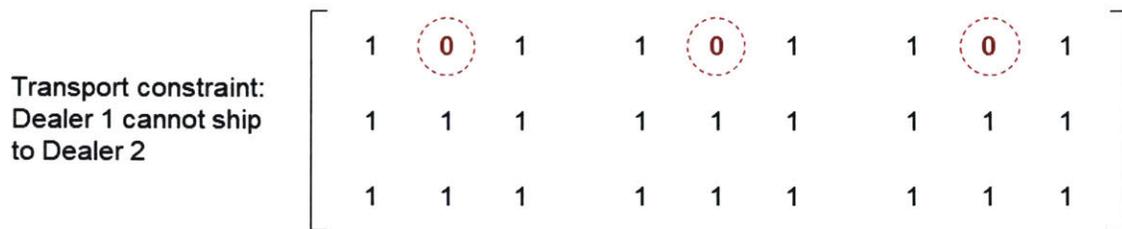


Figure 20: The full set of network edges, A, and the durable transportation constraints, B

In order for a swap edge to be viable, the origin dealer must have a unit of inventory which matches the model and configuration of a backorder at the termination dealer. The presence of this condition is identified by a third matrix, C, with dimensions  $ijk$  which is constructed by taking the row-wise product of an  $ik$  matrix representing origin dealer inventory and a  $jk$  matrix representing termination dealer backorders (Note: the definition of a 'row-wise product' is clarified in Appendix: Row-Wise Product Description). These three matrices have been constructed for the illustrative example above and are shown in Figure 21.

**Matrix C<sub>I</sub> ( available inventory)**

$$\begin{matrix}
 & & & \underbrace{\hspace{10em}} \\
 & & & u \text{ (Model)} \\
 \left. \begin{matrix} s \text{ (Origin} \\ \text{Dealer ID)} \end{matrix} \right\} & & \left[ \begin{array}{ccc} 0 & 1 & 1 \\ 1 & 0 & 1 \\ 1 & 1 & 0 \end{array} \right]
 \end{matrix}$$

**Matrix C<sub>II</sub> ( backorders)**

$$\begin{matrix}
 & & & \underbrace{\hspace{10em}} \\
 & & & u \text{ (Model)} \\
 \left. \begin{matrix} t \\ \text{(Termination} \\ \text{Dealer ID)} \end{matrix} \right\} & & \left[ \begin{array}{ccc} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{array} \right]
 \end{matrix}$$

**Matrix C = C<sub>I</sub> ◦ C<sub>II</sub> (feasible exchanges)**

$$\left[ \begin{array}{ccccccccc} 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \end{array} \right]$$

Figure 21: Matrix C, representing viable inventory/backorder pairings

The elementwise product of these three matrices, A, B, and C yields a matrix D, describing the quiver across which inventory swaps could actually be executed. Given the fact that thousands of unique configurations exist in Caterpillar’s system, the resulting quiver could still be quite extensive. However, if multiple possible units could satisfy a potential swap, the definition of one such machine as ‘optimal’ is not typically a meaningful distinction. Dealers may have multiple potential criteria that they would use to evaluate which unit would be sent to satisfy the swap. Consequently, this allows the complexity of the network to be further reduced by converting it into a simple directed graph. Specifically, if *any* model/configuration along a given arc direction, *st*, is present in the quiver, that arc is considered to be present in the resulting directed graph. The directed graph itself is defined by a two-dimensional matrix E, which is constructed by taking the maximum of each *st* vector in matrix D along the *u* dimension. The feasible swap quiver, simple directed graph, and conversion of matrix D into matrix E are presented in Figures 22 and 23 below.

**Illustrative feasible swap network (quiver and directed graph forms)**

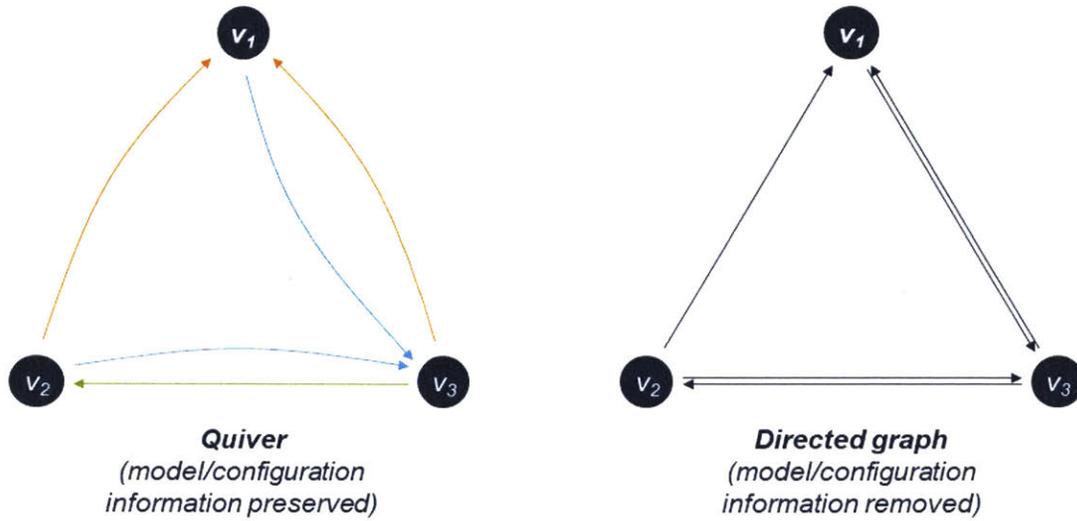


Figure 22: Illustrative feasible swap quiver and directed graph

**Illustrative reduction of Matrix D to Matrix E**

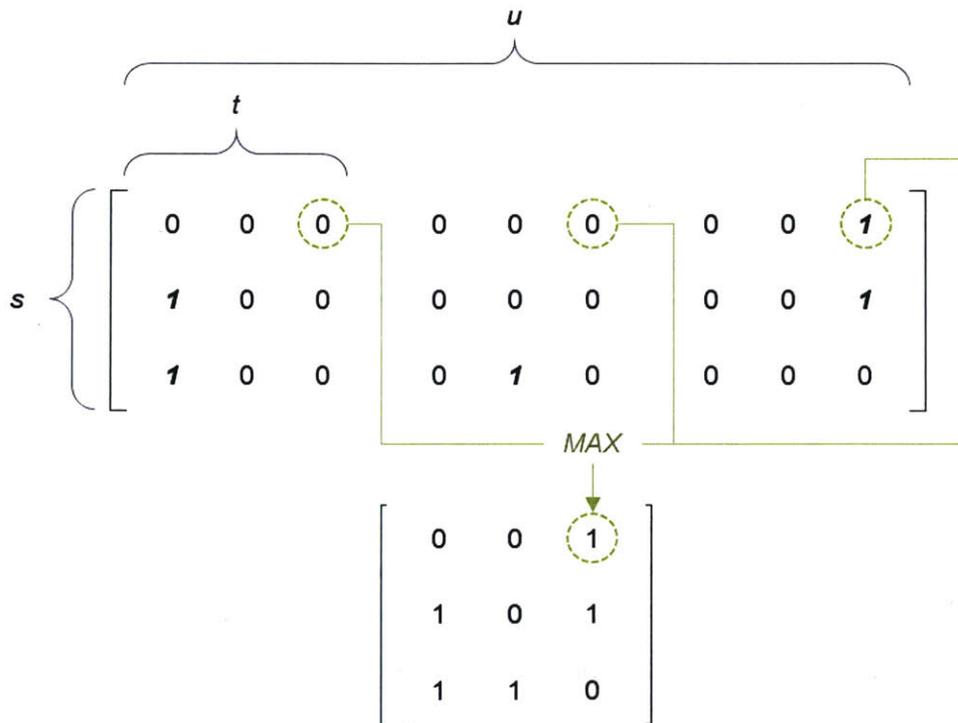


Figure 23: Illustrative reduction of matrix D (feasible swap quiver arcs) to matrix E (feasible swap directed graph arcs)

Using the groundwork above, we can formulate an individual swapping identification optimization problem within the directed graph whose arcs are defined by  $E$  as follows:

$$\begin{aligned} & \max \sum_{(s,t) \in E} x_{st} \\ & \text{Subject to:} \\ & \sum_{(s,t) \in \delta_s^+} x_{st} - \sum_{(s,t) \in \delta_s^-} x_{st} \begin{cases} = 0 & \text{if } s \neq PDC \\ \geq 0 & \text{if } s = PDC \end{cases} \\ & \sum_{(s,t) \in \delta_s^+} x_{st} \leq 1 \\ & \sum_{(s,t) \in \delta_s^+} x_{st} - d_{st} \leq 0 \end{aligned}$$

This optimization problem seeks to find the longest cyclic path (or acyclic, if the PDC is present) present in the directed graph describing the dealers, the durable transportation constraints between them, and the feasible swaps as defined by a pair of dealers with at least one matching inventory / backorder pair. The first constraint is the ‘conservation of mass’ constraint which dictates that a dealer will only give up a unit of inventory in a swap if it also receives a unit of inventory in the same swap path. The PDC, however, may send inventory to dealers without receiving a unit of inventory in return and thus acts as an ‘altruistic donor;’ in this case, the terminal dealer in the chain receives a machine without sending one (but must still pay for the wholesale cost of the machine it receives). The second constraint dictates that the optimization problem may not utilize the same arc more than once. The third constraint accounts for the durable transportation constraints and temporary inventory / backorder match constraints.

Since the full network quiver was reduced to a directed graph and information about machine and configuration was lost, the solved optimization problem merely gives information such as “*Dealer I sends a machine to Dealer II.*” In the illustrative example above, the solution would indicate that Dealer I sent a machine to Dealer III, Dealer III sent a machine to Dealer II, and Dealer II sent a machine to Dealer I. Because of the simplicity of this example, the machines sent are visually obvious (Model 3, Model 2, and Model 1, respectively). In the case that *multiple* machine configurations could satisfy one of the edges of the solution, a simple heuristic is used. Namely, the sending dealer elects to send the machine whose current inventory level is highest relative to that dealer’s target position. The logic behind this heuristic is that it will limit the likelihood that the sending dealer finds itself in a stockout position in the ensuing period as the consequence of its executed swap.

After an optimization solution is processed and the model / configuration for each edge of the swap has been identified, the inventory and backorder list for each dealer is temporarily updated. The full network optimization problem is then set up once again with the single revision that matrix C, which reflects the match between extant inventory and backorders at a pair of dealers, will take into account the swaps executed according to the first optimization. This process is repeated until a stopping condition is met; in this case the stopping condition is:

$$\sum_{(s,t) \in \delta_s^+} x_{st} = 0 \text{ for every } s \neq PDC$$

Namely, if the PDC is the only node participating in *sending* machines as part of a swap, the swap identification is terminated (since this optimization result is identical to the seventh step of the process outlined in Figure 13, it can be more efficiently executed without optimization).

Once the optimization cycle reaches its stopping conditions, the backorders present prior to the optimization are restored, but the inventory at each dealer is updated to reflect the results of all of the identified swaps. In the next period,  $t+1$ , the backorders will be served by the inventory received via swaps. This implicitly assumes that the Caterpillar Last Mile network will require one full month to physically execute swaps; this is likely a conservative assumption.

It should be noted that the directed, longest path optimization problem is NP-complete. Consequently, computation time may be a concern in applications where the optimization needs to be solved with high frequency or the scale of the network is extremely large. The full-complexity (50 dealer, 1-PDC, full machine configuration complexity) simulation discussed in this thesis requires approximately four hours to solve on a typical laptop PC. In practice, the dealer swapping optimization problem will only need to be solved on an approximately weekly cadence, and ADSD-N (which is the largest contiguous block of Caterpillar dealers) only has approximately fifty unique nodes; consequently, this optimization problem will be feasible from a business process perspective. However, the scope of this problem does make a full Monte-Carlo simulation with varying stochastic demand patterns infeasible. Consequently, the results of this thesis are based upon a limited number of runs (seven).

#### 4.2.5. Model Validation Cases

Prior to running test cases with the model, a series of contrived validation cases were defined and run to ensure that the results obtained from the swapping model matched the expected results. These validation cases were conducted on the first generation model prior to dealer syndication; the second generation model was then benchmarked against the first generation model to ensure that it generated the same results under the same baseline conditions. The following table summarizes the scenarios that

were run in initial validation, the objective of each contrived scenario, and the results that were obtained.

Set Up	Objective of Test	Test Results
Run full complexity model (i.e., 13 dealer network with PDC) under 12-month, 24-month, 36-month, and 48 month time spans	Model should produce the same results over the same time period under each run range. Confirming this will allow shorter validation cycle time.	First 12 months of data in all four scenarios is identical; first 24 months in scenarios 2-4 is identical; first 36 months of data in scenarios 3 and 4 is identical
Run dealer only model (13 dealers, no PDC) with durable transportation constraint matrix zeroed out for 12 month time span	Tests that iterative, single period model ("Single Period Swap") produces the same results as the multi-period no swap model ("No Swap Calcs") under the condition of no match network	The results from the swapping and no-swapping models were identical
Run dealer only model with 3 contrived dealers; each dealer only inventories one machine and only has demand for another	Tests that swaps are being correctly identified by network optimization	In each period, the model identifies all expected swap triangles which are set up in this contrived example.
Run model with 12 dealers, 9 dealers, 6 dealers, and 3 dealers, where each dealer has an identical average demand & inventory position (but different demand profiles)	Tests that the number of swaps identified by the model increases as the number of dealers in the network increases.	Tests were run for each integer number of dealers (2 through 13); the number of swaps identified increases slightly super-linearly with respect to the number of dealers included.
Run model with identical dealer set with and without PDC	Tests that presence of universal donor increases the average length and number of matches identified by network optimization	The average number of swaps per period increased from 9.0 to 9.6
Run model with 13 dealer network but only one type of machine inventoried and only one type of machine demanded	Tests that dealers seek to serve their own demand first and must receive a machine to make a swap; no matches should be identified since the dealer will either have enough inventory such that it has not backorders or it will not have any inventory to swap.	No swaps are identified in the matching case; in each month, any given dealer <i>either</i> had residual on hand inventory or backorders, but not both.
Run model with 3, 6, and 9 month lead times (must run model for 24 months to allow transitory period to be worked out of system)	Tests that matching provides a greater relative advantage when factory lead times increase; benefit should increase asymptotically. If factory lead times were one month, there should be no benefit (dealers could get new machines from factory as fast as from another dealer)	Overall sales advantage for the swapping case vs. the no swapping case was 6.8, 13.3, and 17.4, respectively.

Figure 24: Model validation summary

In each of the validation scenarios, the model performed according to the expectations of the contrived scenario. This suggests that the model is indeed identifying and executing swaps according to the assumptions, inventory control methodology, and swap optimization outlined in section 5.1.

#### 4.3. Dealer Syndication

Upon validating the first generation model, a series of dealer syndication meetings were set up and conducted in order to generate feedback from the potential customers for a network swapping system. In-person meetings were conducted with sales operations managers representing five of Caterpillar's ADSD-N marketing districts, and the preliminary model concept and results were also presented at a Dealer Advisory Group (DAG) meeting, which assembles Medium Wheel Loader product managers from a set of high volume dealers. The objective of these meetings was threefold:

- Begin **developing 'customer pull'** from the dealers for a centrally managed network swapping mechanism
- Determine **which aspects of the model need to be refined and where detail must be added** in order to reasonably simulated the commercial behavior of Caterpillar's dealer network
- **Understand existing process, technology, and behavioral barriers** to implementing a network swapping mechanism

Overall dealer reaction to the concept of network swapping was positive. All of the dealers that were consulted expressed interest in the implementation of some type of enhanced swapping mechanism, and only one of the dealers exhibited significant reservations about the proposed network swapping mechanism. In that specific case, the criticism was primarily rooted in a process concern: the dealer felt that a centrally cleared swapping mechanism would result in the receipt of machines that did not actually match the dealer's requirements. The potential reason for such a mismatch is that dealers sometimes make modifications to machines which are not reflected in the database that is visible to Caterpillar. This represents a legitimate concern which could threaten the viability of a network swapping mechanism; a proposed method for resolving this data issue is discussed in the Implementation Pathway section.

The feedback regarding model refinement was also consistent: each dealer identified machine complexity as a critical factor in accurately modeling the behavior of a proposed inventory swapping mechanism. In one case this was made explicitly clear: on the day of the visit to the dealer, the sales operations manager was seeking to source a relatively common machine (a medium-size dozer) equipped with two rare options; at the time of our visit, only two such machines existed in the Caterpillar network. As a consequence of this feedback, machine complexity was prioritized for inclusion

in the second generation model. Several dealers also expressed an interest in expanding the aperture of potential swaps beyond general Construction Industries (CI) machines to smaller Building and Construction Products (BCP) machines with the caveat that multiple BCP machines may need to be traded for a single general CI machine. This extension is not covered in this thesis but is discussed in the Conclusions and Extensions section.

The discussion of existing process, technology, and behavioral barriers yielded some of the most interesting feedback from dealers and was used to inform the proposed Implementation Pathway. Specifically, the following potential challenges were identified during dealer conversations:

- Several dealers expressed concerns that what appears to be **inventory in a centralized system may not actually be tradeable inventory**; machines may be earmarked as demonstration vehicles, rental vehicles, or they may be vehicles which are earmarked for sale to a specific customer. Any centralized swapping system would have to be able to differentiate between these cases.
- There is some concern that dealer may become **hesitant to trade during periods of supply constraint**, despite the fact that it offers an important additional degree of freedom. Dealers may have a psychological resistance to giving up any machine sitting on their lot, even if they will receive another machine in return.
- The current inventory tracking system has some **gaps with regards to in-transit inventory**. Ultimately, it would be valuable if the swapping system could be used to re-route inventory inbound to dealers from the factory rather than simply moving machines from dealer to dealer, but the current inventory tracking system may not be equipped to handle this.
- Several dealers expressed concerns about **potential mismatches between the actual as-is configuration of a machine and the configuration which is being tracked** in the current inventory database (MPIM). This is a particular issue when dealers use kits to make modifications to machines after receipt from the factory.
- Determining the **appropriate cadence for swapping may be challenging**, since it will have to be synchronized across dealers. Although inventory planning is done on a monthly basis, the current ad hoc swapping cadence is much shorter; dealers often need a swap to be reconciled within a week.
- If degrees of freedom are allowed with respect to the configuration of a machine (e.g., a dealer will accept either a standard or deluxe cab), then some cap may have to be set on the allowable wholesale price of the machine that an algorithm identifies for a trade. If a dealer is attempting

to swap for a customer that is price sensitive, there is a **risk of receiving an 'over-spec'd' machine** that has too many expensive options.

- Some dealers expressed **resistance to swap chains with more than two or three dealers** (at least initially) because of the perceived risk that one of the dealers in the chain could renege when the swap is being physically reconciled.

## 5. Simulation Results

As discussed in section 5.2.4, the nature of the optimization problem and scope of the network being considered is such that a full Monte Carlo simulation is infeasible. In order to ensure that the results in this section have not been skewed by an atypical demand profile, seven baseline runs were conducted. In the event that obvious outliers had been observed, the results would have been presented both with and without those runs. However, the baseline results are relatively stable across all seven observed runs, suggesting that these results are robust and would not deviate significantly if a full Monte Carlo simulation were conducted.

### 5.1. Baseline Network Scope

The Baseline scenario was structured to approximate a fully implemented network swapping mechanism with participation from all of Caterpillar's ADSD-N dealers under a 'steady-state' demand condition. Specifically, the Baseline scenario is defined according to the following parameters.

- The network includes all 46 dealers in the United States as well as Caterpillar's PDC, which preferentially behaves as an altruistic donor for the swapping network. The Baseline network does not include the four Canadian dealers; the justification for this is that stakeholder interviews suggest that swapping across international borders may be too complicated and expensive to be feasible.
- Each of Caterpillar's dealers maintain a consistent inventory target of three months of inventory. Although this figure varies dealer to dealer, this is an approximate rule of thumb that most dealers adhere to.
- For all eight sales models, there are three configuration parameters which are considered 'critical' to match: the boom type (none, standard, short-lift, high-lift, and logger), hydraulics (2-valve, 3-valve ready, 3-valve, and 4 valve), and axles (manual locking differential and auto locking differential). This list of critical parameters implies that the other vehicle configuration parameters can either (1) be modified via delayed differentiation or (2) are not sufficiently important to the customer to prevent a sale.
- As a heuristic for preventing swaps with logistics costs that exceed the benefit of the swap, the top quartile of transportation edges (according to cost) have been eliminated. Practically, this means that direct swaps will not occur between dealers in the Seattle and Harrisburg, Seattle and Louisville, and Seattle and Atlanta districts.
- Dealers behave as though they are ignorant of future customer demand beyond average customer behavior. Within the simulation, this means that dealers do not anticipatorily order to

meet any customer orders, and they are willing to swap any unit of on-hand inventory (after a selling period has been completed) to receive a machine which addresses a backorder need.

- Customers who are backordered are willing to wait indefinitely to receive a machine; there is no order attrition or lost orders as a function of a customer waiting to receive a machine.
- Factory lead time is fixed at three months.
- The only machine class considered in the Baseline scenario is Medium Wheel Loaders (MWLs). This would be the case if (1) network swapping were only implemented in the MWL product group or (2) if different machine classes were not swapped for each other.

Once the Baseline has been defined, a number of alternative scenarios will be considered to examine sensitivities and alternative implementations of the network swapping concept.

As noted above, the Baseline network results were simulated against seven simulated demand profiles. The gross characteristics of each demand profile are summarized in Appendix: Baseline Demand Scenarios, Gross Description. One characteristic which should be noted about all of these demand profiles is that they assume that long-term demand is stable. In reality, macro-economic trends may result in long-term demand cycles that are not considered directly in this simulation. In general, however, an 'up' cycle would result in a situation where the value chain becomes constrained and factory lead times would increase, while factory lead times would decrease in a 'down' cycle. The impact of changes in factory lead time are examined in section 6.3.2.

## 5.2. Baseline Results

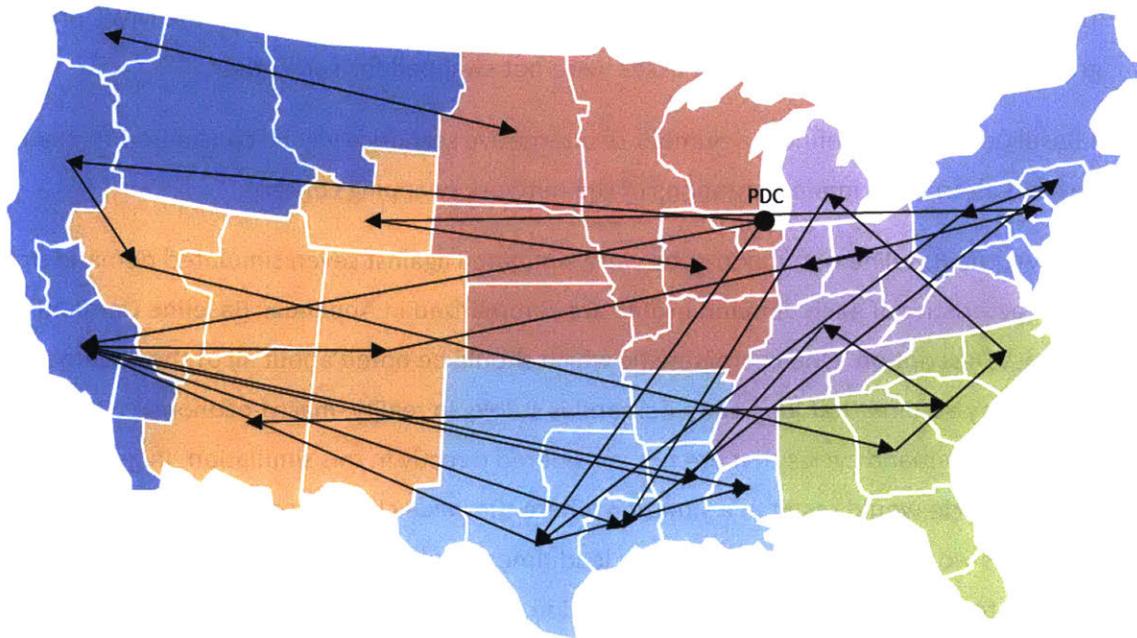
The baseline simulated commercial network displays several results that are characteristic of virtual pooling. Namely, on-hand inventory levels are reduced, customer backorders are reduced, availability (i.e., lead time to customers) is improved, and sales are accelerated. Within the time frame considered, the incremental value of accelerated sales exceeds the discounted logistics cost; the proposed baseline network is NPV positive.

### 5.2.1. Swap Characteristics

The potential benefits and costs derived from network swapping are a direct function of the dealer to dealer trades which are identified and executed. Each trade represents the movement of a unit of inventory from a location where it is accruing inventory holding costs but not immediately being used to a location where it can directly address a backorder. However, each trade also imposes a logistics cost on the dealer network. Characterizing the number and type of swaps identified in a given scenario provides an important indicator of the expected impact of that scenario.

Figure 25 provides an illustrative example of a full set of swaps identified in one period of one run. The information that can be gleaned from this chart is limited, but it does provide a directional sense of the magnitude of swaps which are identified and executed each period in the swapping simulation.

**Swap map, Baseline scenario, Trial 7, period 12**



*Figure 25: Illustrative single period swap map*

Figure 26 presents the monthly average number of swaps and standard deviation in the number of swaps identified in each of the seven runs conducted for the Baseline scenario. Both of these parameters are relatively stable across all seven runs, suggesting that other results (e.g., inventory reduction, backorder reduction) should also be stable.

	Run 1	Run 2	Run 3	Run 4	Run 5	Run 6	Run 7	<b>Average</b>
Mean monthly swaps	42.0	40.2	40.9	41.4	39.0	39.9	40.6	<b>40.6</b>
Standard deviation	8.18	9.26	9.18	8.97	8.98	7.87	7.21	<b>8.52</b>

*Figure 26: Overall swap characteristics, Baseline scenario*

Finally, Figure 27 characterizes the swaps identified in each run according to the type of edge along which the swap is being executed. A PDC edge is one where the PDC is functioning as an altruistic donor to initiate a swap cycle, an intra-district edge is one where two dealers within the same district are swapping inventory, an adjacent edge is one where swaps are being made between two districts which

are geographically contiguous, and a removed edge is one where a swap is being made between two geographically separated districts. This classification provides a short hand proxy for the cost associated with the swaps that are being made, with each of the aforementioned classifications being more expensive than the prior (PDC edges are 'free' regardless of the destination because this fee is already absorbed into the wholesale cost of the machine). It should be noted that the optimization algorithm is 'blind' to the cost of the edges that are utilized to execute swaps insofar as that cost is not incorporated into the objective function; this methodology is critiqued in the Conclusion.

**Swap class breakdown, Baseline scenario**

Total swaps identified, 48 month period (Normalized, Average total swaps = 100)

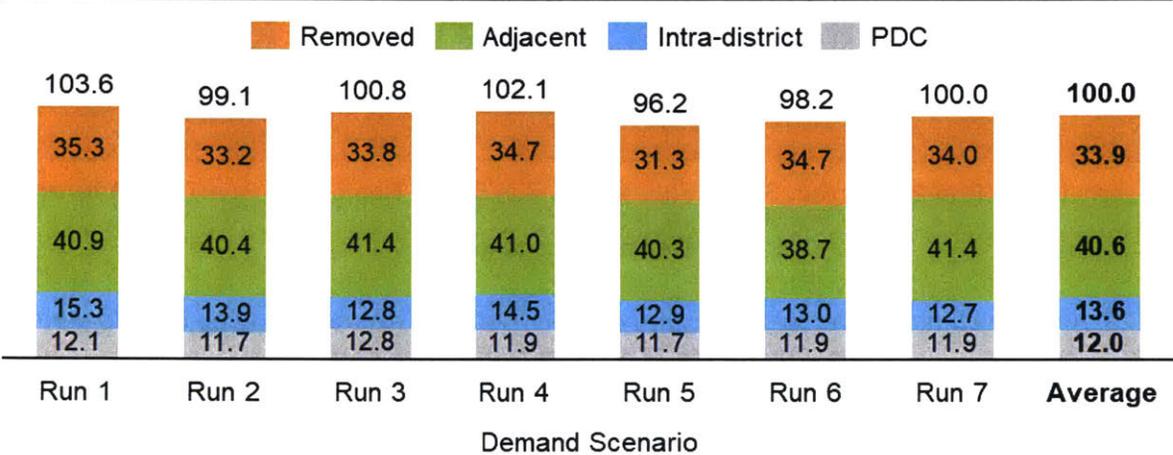


Figure 27: Total swaps classified by edge distance, Baseline scenario

5.2.2. Inventory Impact

In each of the seven observed Baseline runs, the overall on-hand inventory in the Caterpillar dealer network was reduced significantly. Disguised on-hand inventory results are presented in Figure 28; this inventory is the average of period beginning inventory (after factory, PDC, and swap arrivals) and post-sale inventory.

### Average on-hand inventory, Baseline scenario

# of MWLs (normalized to average no-swap inventory across all Baseline runs)

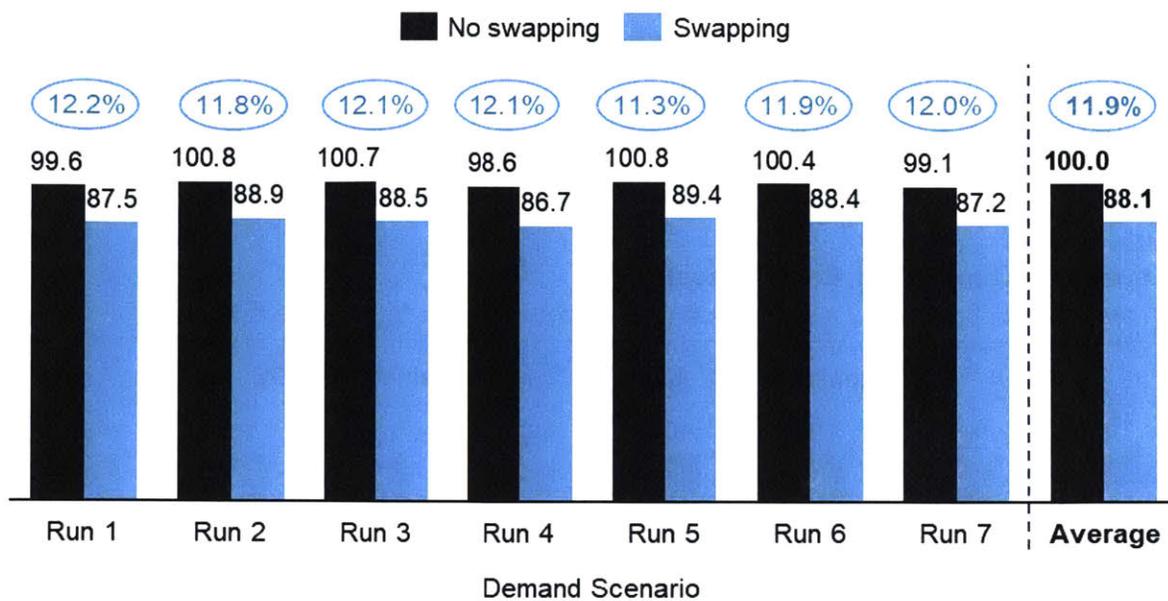


Figure 28: Average on-hand inventory, Baseline scenario (Note: the average no-swap inventory has been normalized to 1,000 units to disguise data)

It should be noted that this inventory reduction is not the function of a conscious inventory policy change made in response to the implementation of network swapping; each dealer is still ordering to the same target inventory position. Rather, this observed inventory reduction is a function of the fact that inventory is moving through the system more quickly – the latency from inventory receipt to sale (or swap) is being reduced.

The on-hand inventory as a function of time during the simulation for a single run is presented in Figure 29. The overall reduction in inventory is relatively stable across the observed time frame. The transitory response observed in the first two months is a function of the fact that each dealer is assumed to start the simulation with its full inventory position on-hand (i.e., pipeline stock and backorders both initially start at zero).

### On-hand inventory, Baseline scenario, Run 7

# of MWLs (normalized to average no-swap inventory across all Baseline runs)

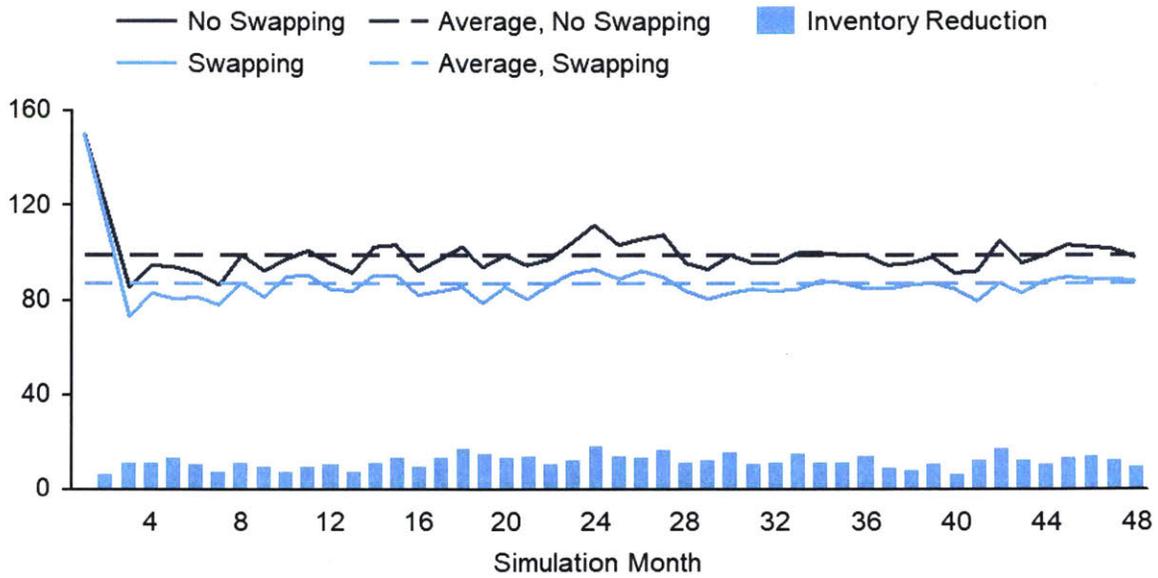


Figure 29: Single run on-hand inventory, swap vs. no-swap

Figure 30 shows that the on-hand inventory reduction is heterogeneous across dealers. In run 7, for example, the minimum inventory reduction observed was 0.0% while the maximum was 21.1%. This heterogeneity was not an obvious function of dealer size – both large and small dealers were present at both ends of this range.

### Dealer level inventory reduction, Baseline scenario, Run 7

Dealer count vs. level of inventory reduction (n = 46)

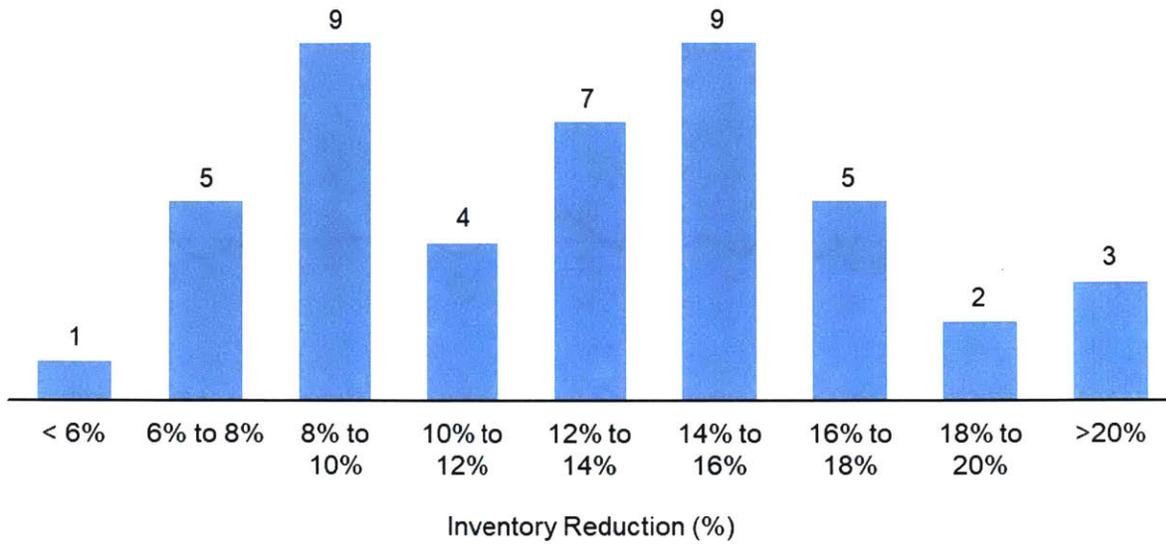


Figure 30: Histogram of dealer count versus inventory reduction level

#### 5.2.3. Backorder and Availability Impact

The second phenomenon which should be observed in an initiative which approximates virtual pooling is a reduction in customer backorders. The observed reduction in customer backorders across the seven baseline runs is presented in Figure 31; the average reduction is 17.5%.

### Average customer backorders, Baseline scenario

# of MWLs (normalized to average no-swap backorders across all Baseline runs)

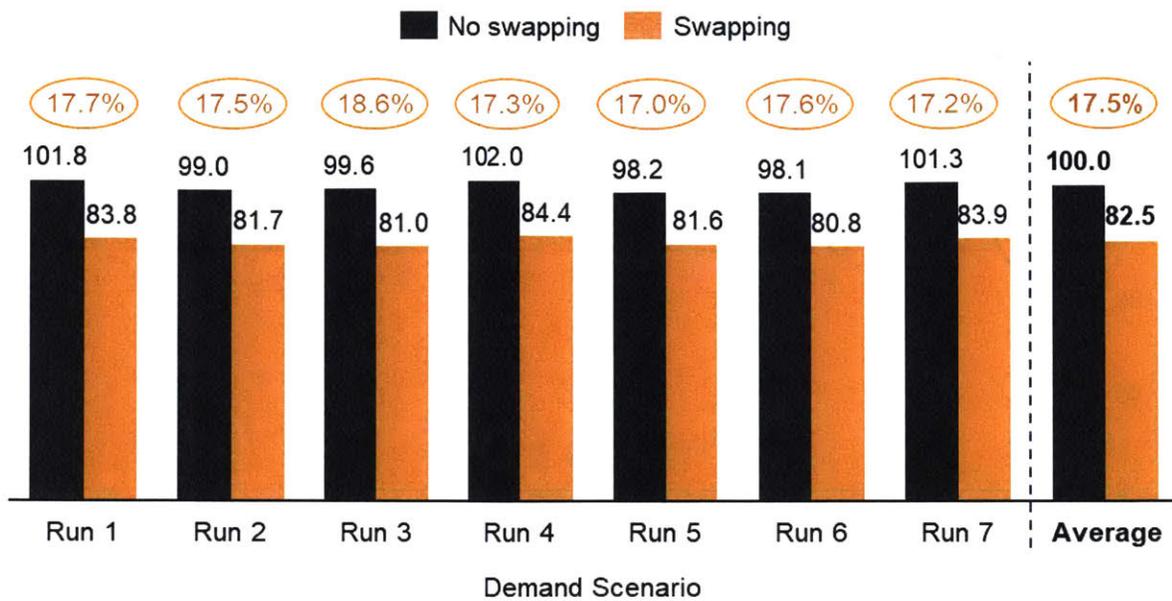


Figure 31: Average customer backorders, Baseline scenario

Observing how backorders evolve over time for a single run provides additional information about how relative customer backorders evolve over time; run 7 is presented in Figure 32. Once again, a transitory effect is present as the dealer’s inventory position reaches its steady-state mix of on-hand and pipeline inventory. It’s notable that the reduction in customer backorders realized by swapping is correlated to the baseline level of backorders itself; that is, the swapping mechanism provides a greater reduction in backorders when the overall level of backorders is higher. This result can be seen more easily in the scatter plot presented in Figure 33. This result suggests that the efficacy of network swapping coincides with the periods where customer demand increases and the supply chain becomes constrained.

**Backorders, Baseline scenario, Run 7**

# of MWLs (normalized to average no-swap backorders across all Baseline runs)

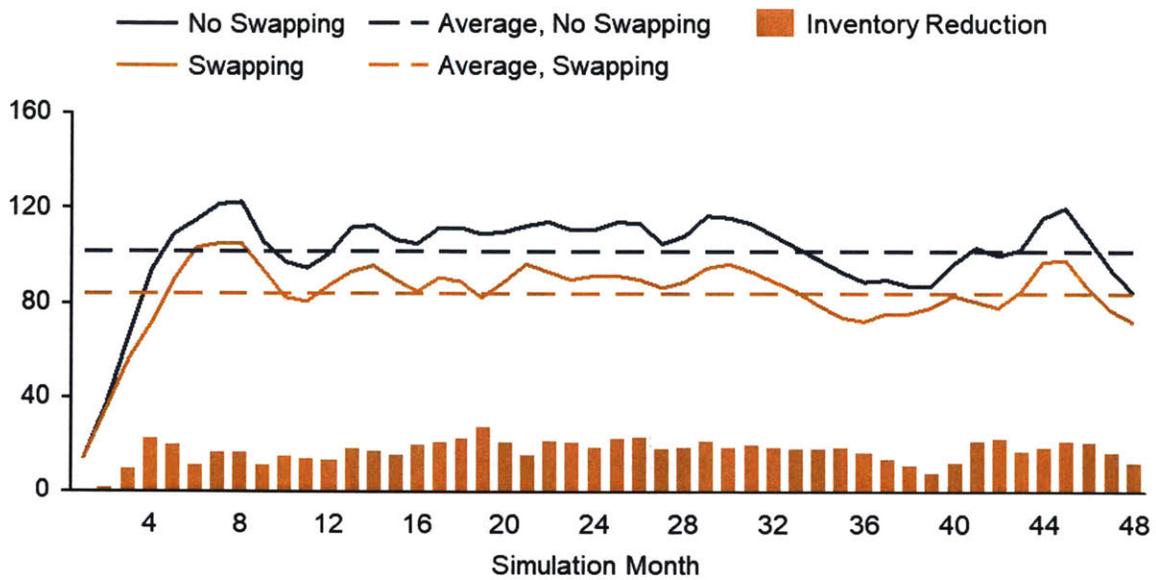


Figure 32: Single run backorders, swap vs. no-swap

**Normalized backorder reduction vs. normalized no-swap backorders, Baseline runs**

# of MWLs (normalized to average no-swap backorders across all Baseline runs)

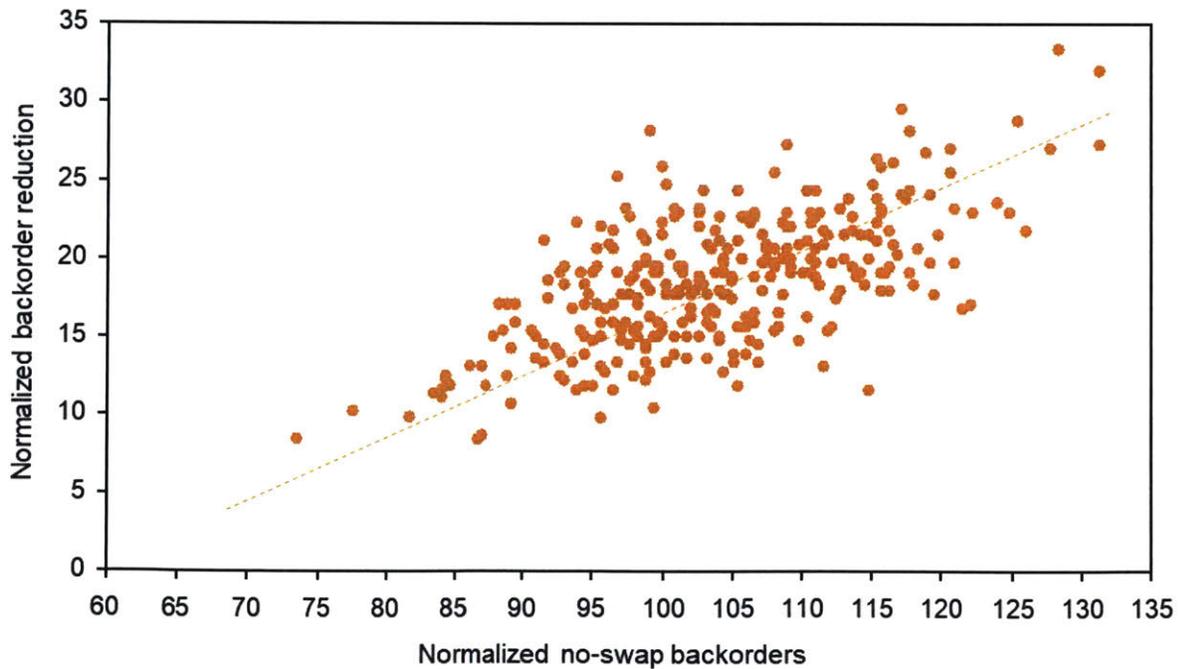


Figure 33: Normalized backorder reduction vs. normalized no-swap backorders, Baseline scenario

One of the metrics that Caterpillar uses to measure the health of its supply chain is Availability, which is defined as the average number of days required to deliver a machine to a customer after receiving an order. In this case, we can apply Little’s Law according to the equations below in order to observe how swapping affects customer perceived availability.

$$L = \lambda W$$

Where L is the number of customers in a system,  $\lambda$  is the arrival rate of customers, and W is the average amount of time spent in the system. Applying this to customer backorders in the dealer network, we can write:

$$\text{Average backorders} = (\text{Customer order arrivals})(\text{Availability})$$

$$\text{Availability} = \frac{\text{Average backorders}}{\text{Customer order arrival}}$$

Figure 34 presents the average availability calculated for each run under the swap and no-swap condition for the Baseline scenario.

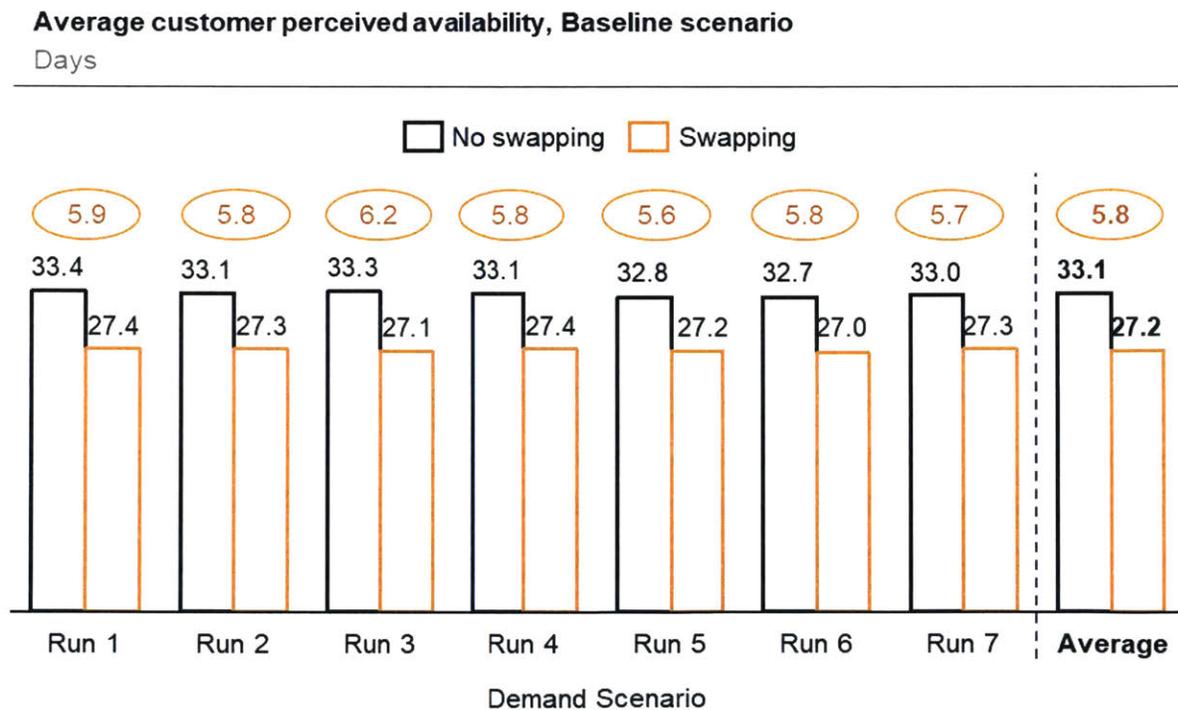


Figure 34: Customer perceived availability, Baseline scenario

This result – an average availability improvement of approximately six days – is significant; the implementation of the network swapping mechanism reduces the average amount of time that a

customer must wait to receive a machine by approximately one week. Such a reduction in availability is a competitive advantage which could be leveraged into increased sales.

#### 5.2.4. Sales Acceleration

From the perspective of Caterpillar’s dealer network, the reduction in customer backorders and improvement in availability manifests as an acceleration in sales and revenue. The absolute acceleration in sales can be observed through the cumulative sales advantage, defined as the cumulative sales under the swapping condition less the cumulative sales under the no-swapping condition. The cumulative sales advantage for a single run of the Baseline scenario is presented in Figure 35.

#### Cumulative sales advantage, Baseline scenario, Run 7

# of MWLs

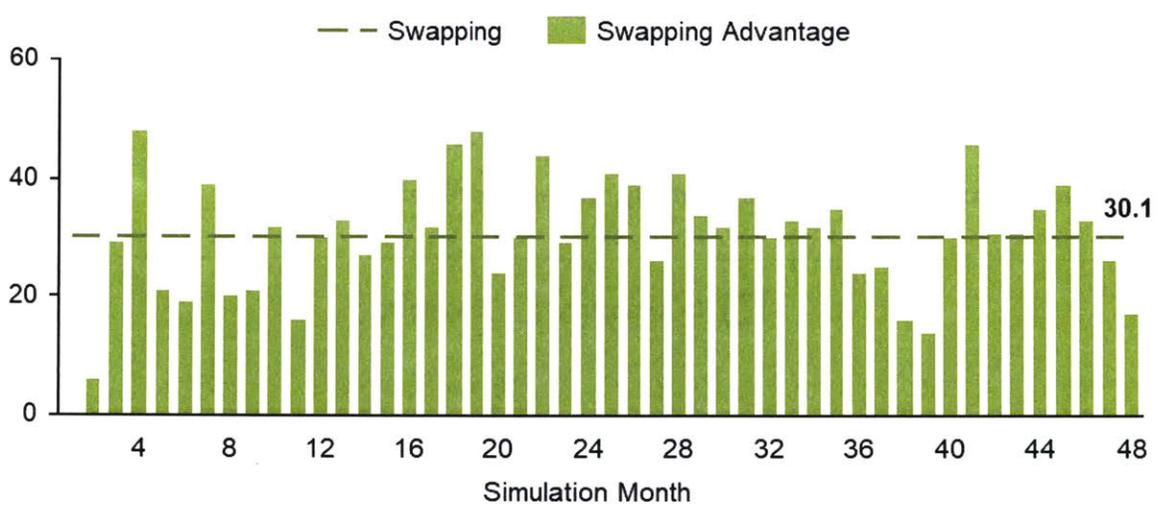


Figure 35: Single run cumulative sales advantage, swap vs. no-swap

In this run, the sales advantage fluctuates depending on the number of swaps identified in the prior period, but the average sustained sales advantage is 30.1 machines across the Caterpillar network. The average sustained sales advantage across all seven runs is shown in Figure 36.

### Average cumulative sales advantage, Baseline scenario

# of MWLs

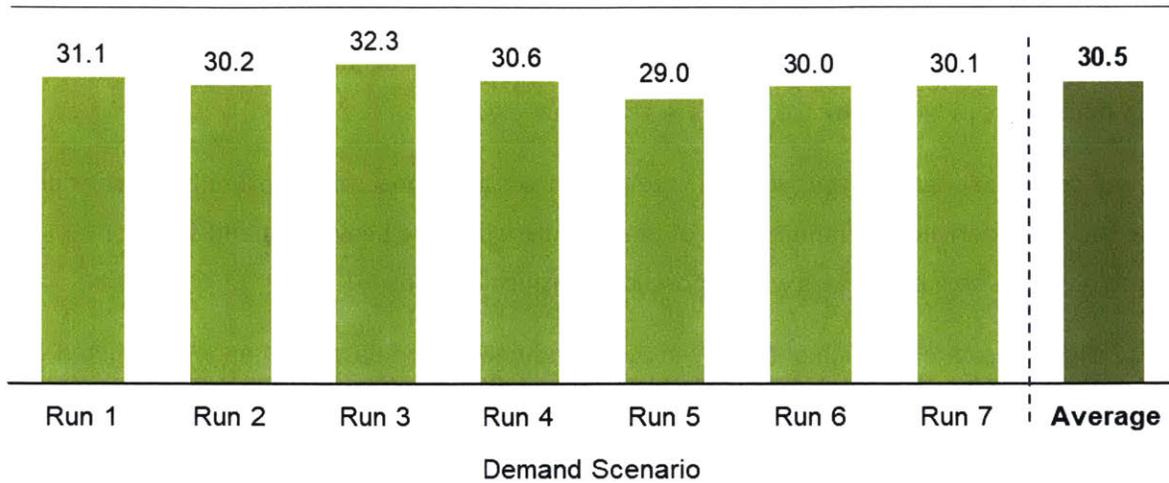


Figure 36: Cumulative sales advantage, baseline scenario

The net impact of this sales advantage is that, as long as the swapping network is maintained, Caterpillar's dealers have effectively captured 30.5 incremental machine sales despite the fact that the underlying demand signal is unchanged.

#### 5.2.5. Network Financial Impact

Ultimately, the objective of implementing a network swapping mechanism is to improve financial returns for Caterpillar and its dealers. There are three potential sources of a cash flow improvement:

- A **one-time increase in sales derived from the cumulative sales advantage described in section 6.2.4**. As long as network swapping is maintained, this sales advantage will persist. Since factory purchases are not accelerated in order to achieve this sales acceleration, this one-time increase is fully attributable to cash flow improvement.
- An **ongoing reduction in inventory carrying cost**. This reduction is composed of two components: (1) the financial benefit from having less capital tied up in inventory and (2) the operating cost benefit from a reduction in firm inventory carrying costs (e.g., the cost to exercise machines that are held in on-hand inventory).
- An **ongoing improvement in revenue derived from additional sales capture**. Voice of customer data suggests that an improvement in availability *does* result in the incremental capture of sales that would otherwise be lost to competitors.

The one-time increase in sales and the reduction in financial inventory holding costs are included in this analysis. For the analysis below, the assumed inventory holding cost is 3.0% per annum; this figure was

selected to provide a conservative estimate of potential financial benefits and to align with the figures reported by dealers themselves.

Firm inventory cost reductions are not considered because they will be idiosyncratic to how each dealer elects to manage its inventory; these costs are also relatively minor.

Finally, the potential ongoing improvement in revenue from additional sales capture is not considered because the value derived from this source of financial benefit is too heavily dependent upon assumptions about how customers will respond to improved availability.

Figure 37 shows the NPV for each of the seven runs conducted in the Baseline scenario, adjusted with a 6-month trailing average (i.e., the NPVs shown are the average of the NPV as of 43 months, 44 months, 45 months, 46 months, 47 months, and 48 months). The reason for this trailing average is that the month to month NPV is somewhat volatile due to the varying sales acceleration observed in each month; the trailing average dampens this effect.

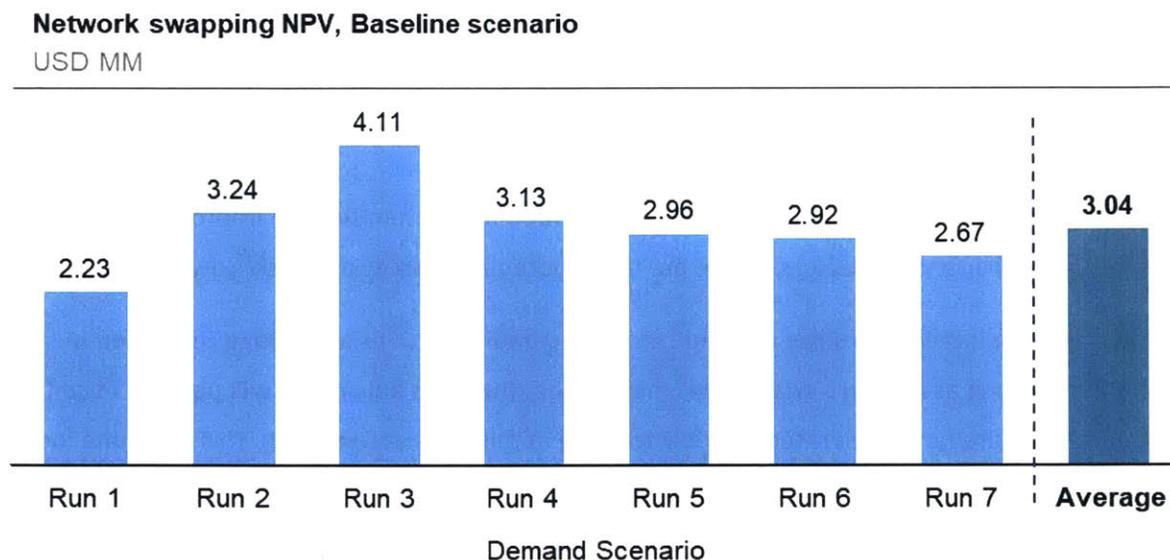


Figure 37: Network swapping NPV, Baseline scenario

The average NPV across all seven scenarios is USD 3.04MM. A simple check suggests that this is reasonable: the total value of accelerated sales (~30 machines at ~\$300K each) is approximately USD \$9 MM while the discounted cost of logistics is approximately USD \$6 MM. At a holding cost rate of 3%, the impact of financial inventory holding cost reductions are relatively small.

Examining the incremental cash flows and NPV for a single scenario provides more detail on the composition of the improvement in dealer network NPV. Figure 38 shows the incremental undiscounted cash flows for run 7, while Figure 39 shows the evolution of NPV, including the 6-month lagging NPV.

**Cumulative incremental cash flows, Undiscounted, Baseline scenario, Run 7**

USD MM

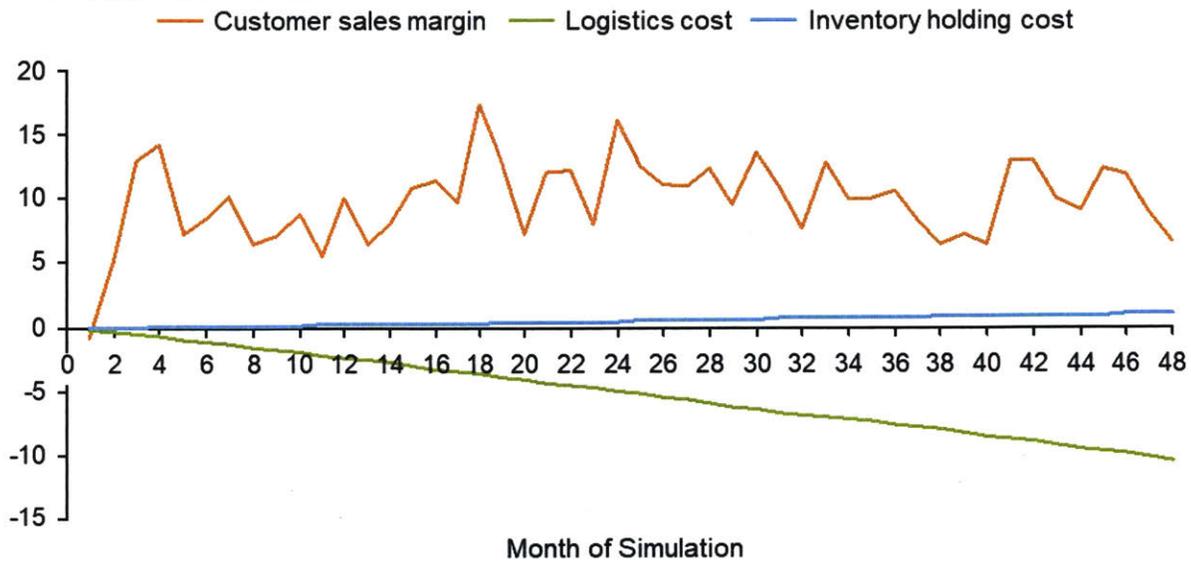


Figure 38: Single run cumulative, undiscounted incremental cash flows from network swapping

**Net Present Value, Baseline scenario, Run 7**

USD MM

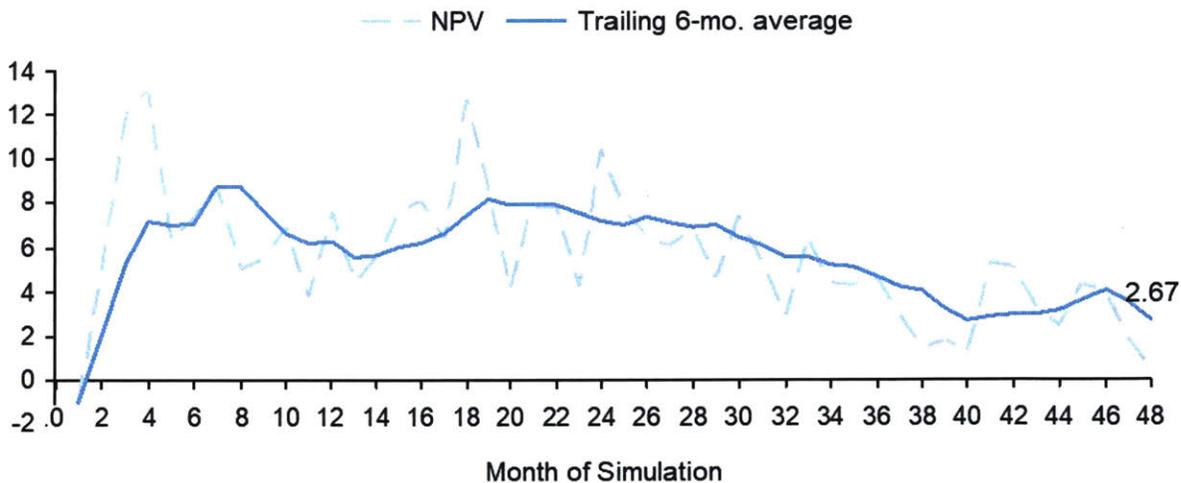


Figure 39: Single run Net Present Value of network swapping

In Figure 39, it's noteworthy that the considered financial *benefits* of accelerated sales accrue almost immediately as the one-time sales gain arrives as soon as swapping reaches its steady state density. Thereafter, the benefit from this source fluctuates around a mean value as period to period swap identification fluctuates. The additional financial *costs*, however, accrue over time as the logistics costs required to sustain the network swapping mechanism are incurred on an ongoing basis. The net effect

of this time differential can be seen in the NPV chart of Figure 39, which shows an NPV which is initially higher and decays over time as more logistics costs are accrued.

Figure 40 shows a histogram of dealer-level NPV improvement for a single Baseline scenario run. As with inventory reductions, there is a significant amount of variance in the financial benefit realized by individual dealers, suggesting that some methodology for pooling these benefits may have to be identified.

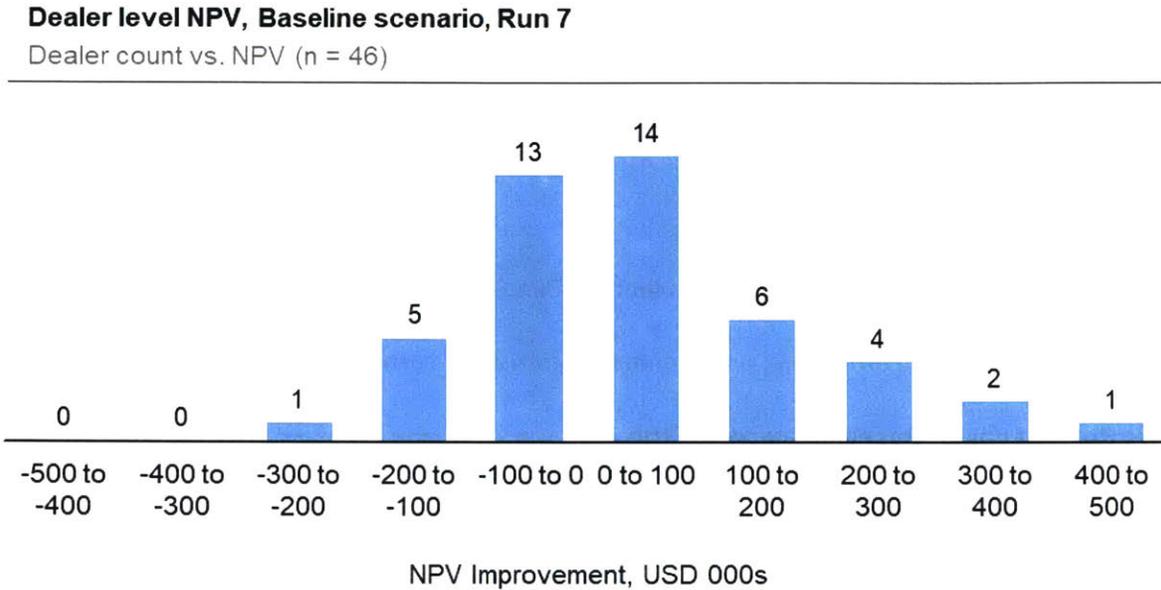


Figure 40: Histogram of dealer level NPV improvement

### 5.3. Network Parameter Sensitivity

This section discusses a number of alternative scenarios which were run to understand how network swapping performs under alternative conditions. In each case, the sensitivities were analyzed with respect to a single base case run (run 7).

#### 5.3.1. Network Scale Effects

In order to observe the effects of network scale, a scenario was modeled in which dealers are circumscribed to swap with other dealers in their own marketing district (see Appendix: Dealer District Map). This scenario is also likely indicative of what a network swapping concept may look like at a relatively early stage of implementation; in order to build confidence in a swapping mechanism, swaps could initially be conducted among a constrained network of dealers with high familiarity.

The impact on swapping density and parameters such as inventory reduction and backorder improvement aligns with the results we may expect. Since the overall optimization problem is effectively

becoming more constrained (edges of the network war being removed), the swap density decreases by approximately 50%, as shown in Figure 41.

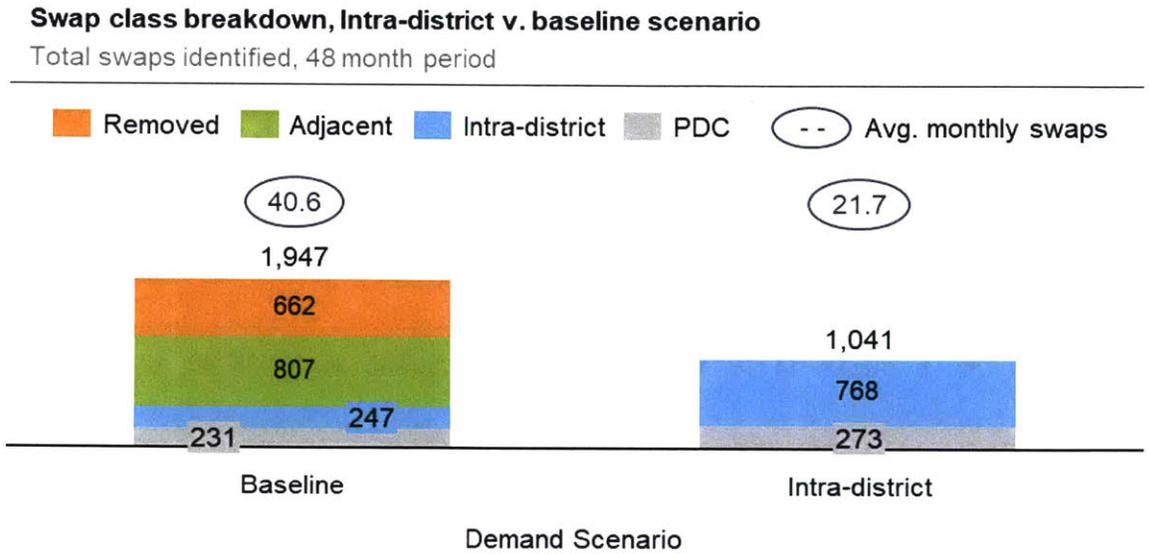


Figure 41: Swap classification, baseline vs. intra-district scenario

Furthermore, network level on-hand inventory reduction decreases from 12.0% in the base case to 5.2% in the intra-district case, and backorder (and availability) improvement decreases from 17.2% to 7.9%. Average sales acceleration drops from 30.1 units to 13.9 units.

Interestingly, however, the overall Net Present Value of the swapping network remains effectively unchanged, with the final 6-month lagging NPV improving from USD 2.67MM to USD 2.87MM, as shown in Figure 42.

### Net Present Value, Intra-district v. baseline scenario

USD MM

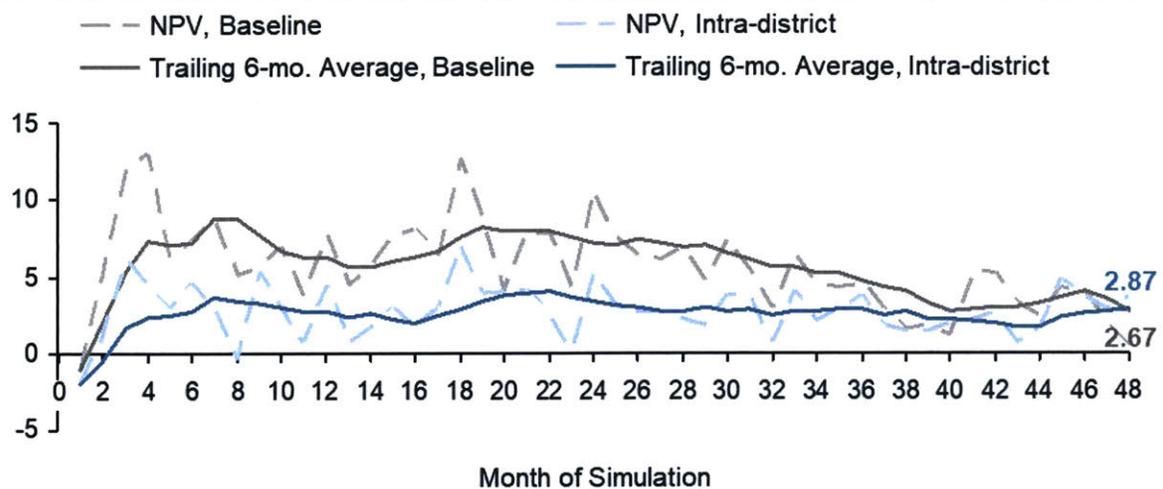


Figure 42: NPV comparison, base case v. intra-district case

The reason for the constancy in NPV despite the reduction in swapping and sales acceleration becomes evident upon observing the cash flow breakdown of the intra-district case and comparing it to the base case; this comparison can be seen in Figure 43.

While the cash flow improvement from sales acceleration drops from an average of USD 9.87MM to USD 3.42MM, the cumulative undiscounted cost of transport drops from USD 10.48MM to USD 1.77MM. The implication of this result is important: in the base case network, some long-edge swaps which are being conducted are destroying value. Furthermore, this suggests that the simplistic longest-path optimization applied in this thesis is likely too reductive for a real network swapping mechanism, and a refined objective function should somehow account for incremental sales and cost. This critique is addressed further in Section 8, Conclusions and Extensions.

### Cumulative incremental cash flows, Undiscounted, Intra-district scenario

USD MM



Figure 43: Undiscounted cash flow breakdown, intra-district scenario

#### 5.3.2. Factory Lead Time Effects

Although factory lead time was assumed to be constant in this simulation, actual factory lead times dilate and contract during periods of sustained demand uptick or drop. Consequently, Baseline scenario 7 was run with 2-month, 3-month (base case), 4-month, 5-month, and 6-month lead times. A 1-month scenario was not observed; the logic of the swapping simulation would prevent swaps from occurring in such a case since swaps are only identified if there is lead time advantage versus the factory order.

In these scenarios, the target dealer inventory position is not adjusted. In periods of sustained factory lead time dilation, dealers would be expected to increase their inventory target in order to avoid a major impact on customer service level (conversely, dealers would likely reduce held inventory if factory lead time grew shorted over a sustained period of time). However, the assumption of a constant inventory position is likely valid for transition periods since there is some latency in how dealers and Caterpillar are able to adjust to macroscopic demand trends. Consequently, these scenarios provide some insight into how a swapping mechanism performs during periods of demand ramp-up or ramp-down.

Overall, swap density is highest in the Baseline (3-month lead time) case (see Figure 44 below). In the 2-month lead time case, dealers are able to serve enough demand from their own on-hand inventory to reduce the frequency with which they must seek swaps. Conversely, as lead time extends beyond 3-months the inventory that dealers have to make swaps declines rapidly, inhibiting swaps even as backorders increase.

### Average monthly swaps identified, factory lead time variation

# of swaps

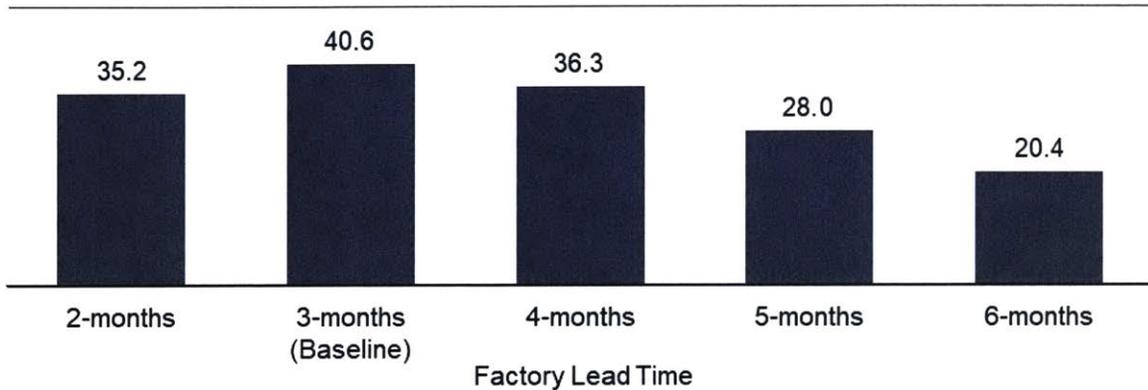


Figure 44: Average monthly swap identification across various factory lead times

Absolute inventory reduction, availability improvement, and sales acceleration are all observed at a 4-month lead time. Proportionally, however, inventory reduction improves monotonically as the average on-hand inventory in the network declines as lead times dilate. Furthermore, the NPV uplift observed also increased monotonically across the lead time range observed; the NPV uplift in each scenario is presented in Figure 45. The justification for this is relatively straightforward: when the factory lead-time is only two months, a successfully identified swap only accelerates a sale by one month; when factory lead-time climbs to six months, a successfully identified swap accelerates the sale by five months. This greater acceleration translates into a greater acceleration of cash flows and – consequently – a greater improvement in NPV for each unit that is swapped.

### Network swapping NPV, factory lead time variation

USD MM

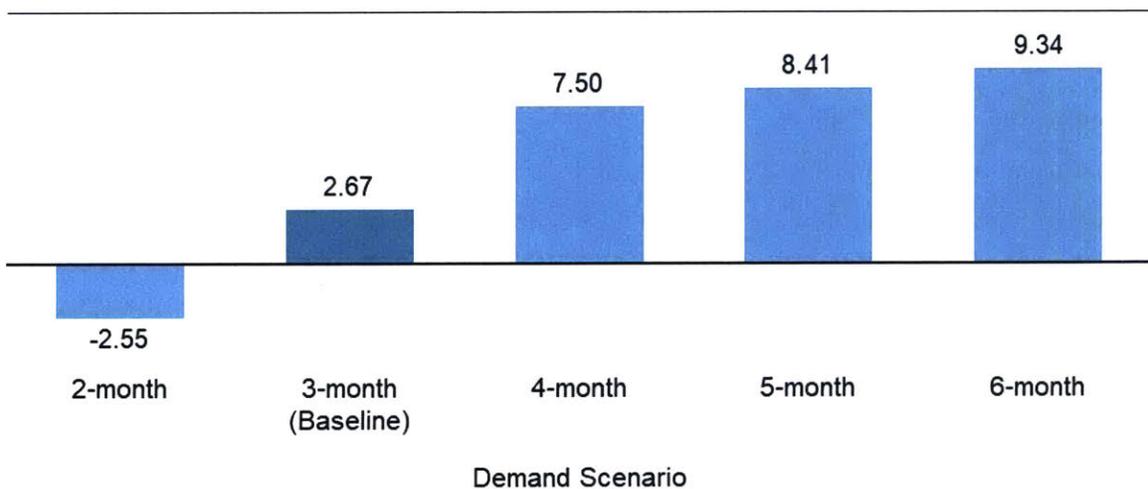


Figure 45: Network NPV impact across various factory lead times

Once again, the negative NPV observed in the two-month lead time case suggests that a significant portion of the swaps being conducted within this time frame are destroying value. This result is somewhat intuitive – since the relative sale acceleration and commensurate financial benefit is declining, only short-distance, low-cost swaps will make financial sense.

### 5.3.3. Complexity Requirement Effects

The final sensitivity which was observed is how the swapping network responds to increasing degrees of dealer configuration specificity. In the Baseline case, only three configuration parameters – the loader boom type, the hydraulic articulation, and the type of differential – were considered. This scenario represents a situation in which Caterpillar has dramatically reduced the configuration complexity of its offering or implemented a very high degree of delayed differentiation. In reality – especially in early implementation – dealers would likely specify their desired machines along a greater number of dimensions. To examine the impact of a higher degree of configuration specificity, 6-parameter and 9-parameter cases were examined. The feature categories specified in each case are presented in Figure 46.

<b>(3-Parameter) Baseline</b>	<b>6-Parameter</b>	<b>9-Parameter</b>
Boom	Boom	Boom
Hydraulics	Hydraulics	Hydraulics
Differential	Differential	Differential
--	Cooling Core	Cooling Core
--	Cooling Fan	Cooling Fan
--	Axles	Axles
--	--	Seat
--	--	Steering Control
--	--	Implement Control

Figure 46: Configuration parameters considered in 3-parameter, 6-parameter, and 9-parameter cases

The result which is most immediately apparent is that NPV improves as configuration specificity increases; this can be seen in Figure 47.

### Network swapping NPV, configuration specificity variation

USD MM

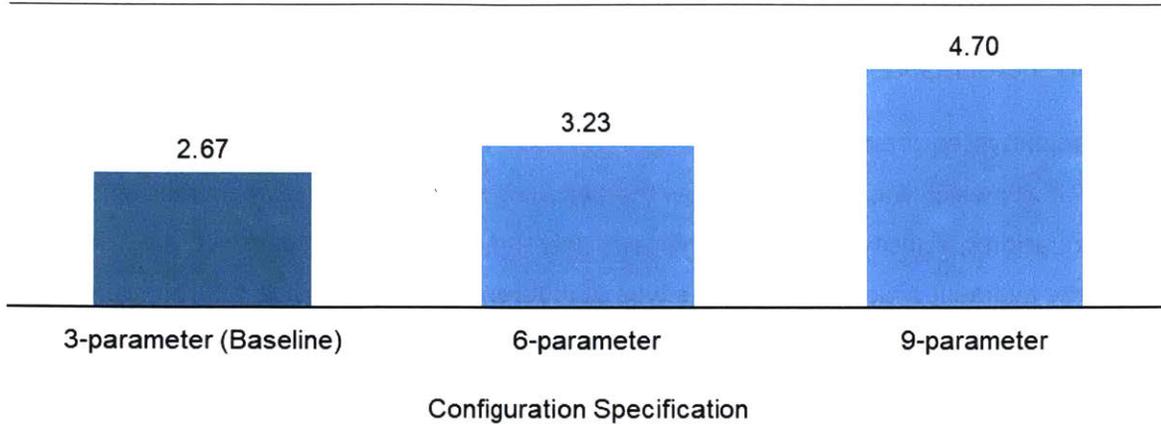


Figure 47: Network swapping NPV across various configuration specificities

This result is counter-intuitive: although optimization constraints are being made more restrictive as the number of parameters specified increases, the network is performing better from a financial standpoint. The source of this unexpected result can be observed through the breakdown of incremental cash flows in each scenario (Figure 48): although the incremental cash flows from sales acceleration (customer sales margin) declines somewhat as the configuration specificity increases, the cumulative logistics cost decreases dramatically.

### Cumulative undiscounted cash flows by source, configuration specificity variation

USD MM

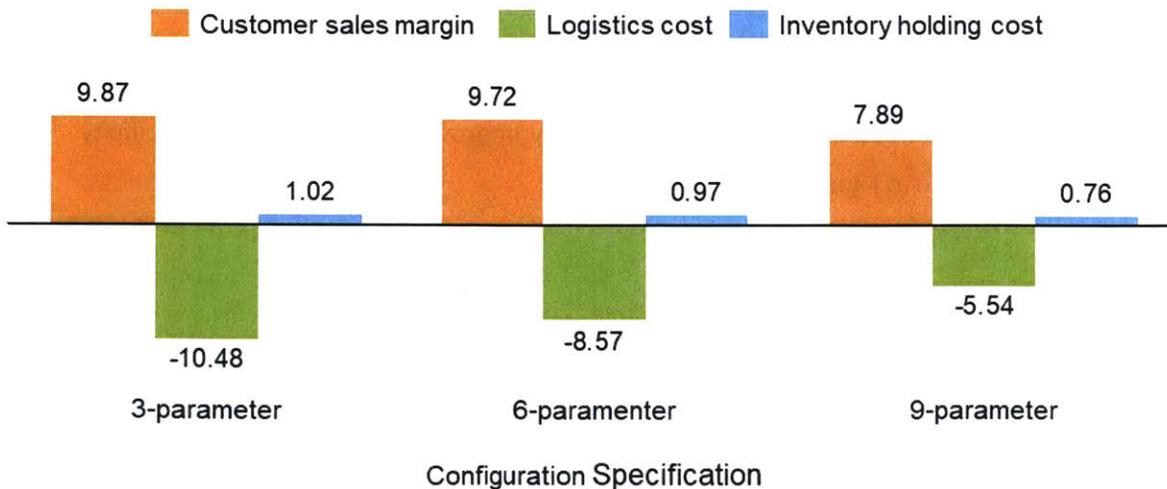


Figure 48: Undiscounted cash flows by source across different levels of configuration specificity

One likely explanation for this effect is that higher configuration specificity implicitly favors swaps with dealers that are geographically close to each other and incur lower logistics costs. Consider, for example, the case of the Cooling Core parameter. The possible options for this parameter are a 6 Fins Per Inch (FPI) or 9 FPI radiator. 6 FPI radiators are preferred in applications where fine airborne debris is high (e.g. logging) or icing is likely, while 9 FPI radiators are strongly preferred in hotter regions. It makes sense, then, that dealers who prefer one configuration or the other would be geographically clustered. Once again, one of the primary implications of this result is that the most effective network swapping mechanism would have an objective function that accounts for the ‘value’ of a swap rather than simply maximizing swap identification.

The other results in the increasing configuration specificity case are more intuitive – swap identification does decrease as configuration specificity increases. Consequently, on-hand inventory reduction, customer availability improvement, and sales acceleration all decline as configuration specificity increases. The results for each of these operating indicators are summarized in Figure 49.

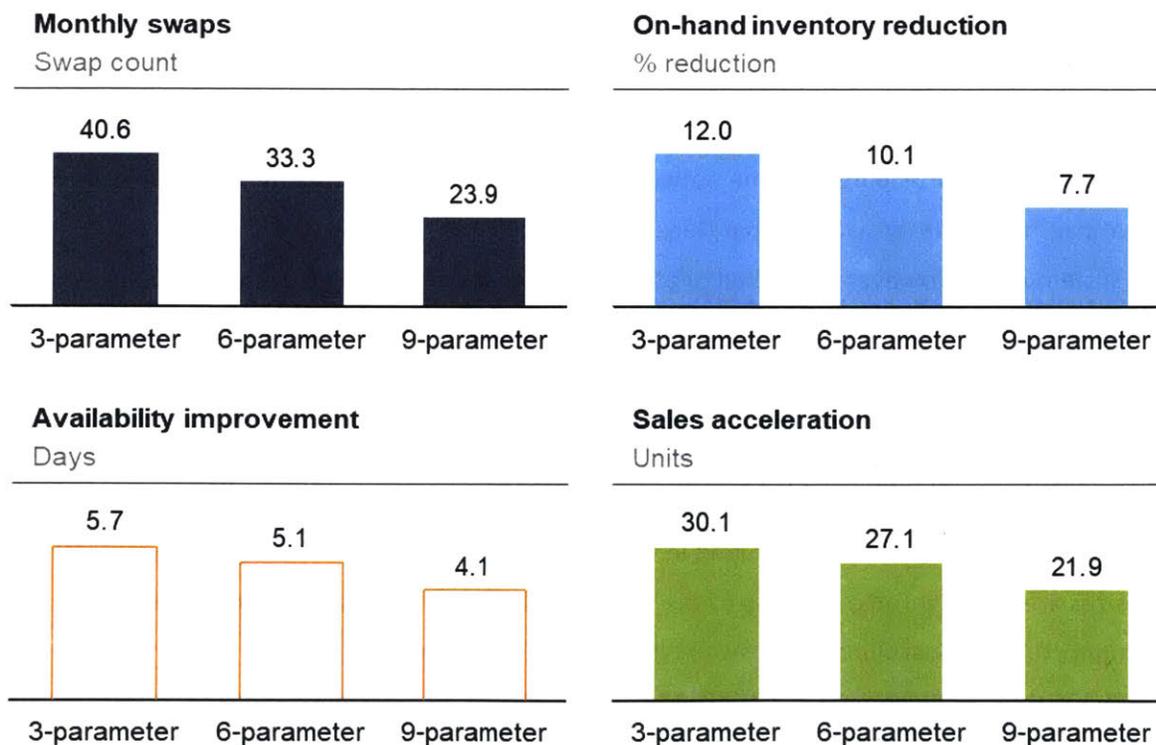


Figure 49: Summarized impact of configuration specificity on operating indicators

## 6. Implementation Pathway

The objective of this chapter is to sketch an overall implementation pathway for the swapping mechanism described in this paper and to highlight some of the barriers that must be overcome to realize a stable, robust swapping system.

### 6.1. Swapping Pilot

The dealer conversations described in Section 4.1 suggest that there is enthusiasm and a latent need for an improved swapping mechanism; however, a carefully structured pilot could provide tangible evidence of real-world impacts for a centralized swapping mechanism. Identifying and executing swaps that (1) accelerate sales and (2) otherwise would not have occurred will provide clear case examples which can be leveraged to justify a full scale implementation.

#### 6.1.1. Pilot Scope

There are three dimensions along which the scope boundaries of an initial pilot must be set: dealer network, product groups, and timeline.

One of the foundations for a successful swapping network is trust; this is also one of the sources of concern that arise during dealer syndication conversations. If dealers do not trust a swapping network, they will elect not to participate, and the network will not achieve the critical mass it needs to yield regular, impactful inventory swaps. Ultimately, this trust can and should be directed towards the central clearing system itself. However, the initial pilot should leverage inter-dealer relationships that already exist. Consequently, the **initial pilot should be circumscribed to a single dealer district**. When selecting which district to pilot, the following dimensions should be considered:

- The **relative dealer size within the district should be examined**; more favorable districts are ones in which the dealers have approximately similar sales volumes, particularly with respect to the products which will be included in the pilot program. Districts with highly asymmetric dealer sizes are less favorable because of the likelihood that very small dealers may not have the opportunity to participate in a swap within the defined timeline.
- The **extent to which dealer districts currently conduct ad hoc swaps** should be weighed. Dealer districts with a higher frequency of swapping are more favorable; these districts will be more familiar with the processes and logistics required to process swaps once they have been identified.
- Finally, the **pilot dealer district should be relatively geographically compact**. This will help ensure that the swapping benefits which accrue during a pilot will not be overwhelmed by the logistics costs associated with executing the swaps.

Given these requirements, it's likely that one of the five eastern U.S. districts – Harrisburg, Louisville, Atlanta, Minneapolis, or Houston – will be appropriate for a pilot program.

The product scope of the initial swapping program should be tractable, but it should also include enough machine volume such that a number of swaps can be identified within the pilot program timeframe. Although the results of this thesis suggest that positive results can be obtained even when the swapping is circumscribed to a single product group (MWLs), expanding these boundaries would increase the likely success of a pilot. Preliminary conversations have been conducted with **Medium-Wheel Loaders, Excavators, and Medium Tracked Product** (primarily dozers). Including these three product groups in the pilot will capture a large portion of the general Construction Industries segment while still maintaining a reasonable scope for manual pilot processes.

The timeline for a pilot program is to some extent mutable. However, the program should be long enough to allow most dealers in the network to be exposed to some swapping without unnecessarily delaying program implementation if the pilot proves to be successful. It's likely that a **three- to six-month pilot** should be sufficient to judge the impact of the pilot swapping network.

#### 6.1.2. Pilot Logistics

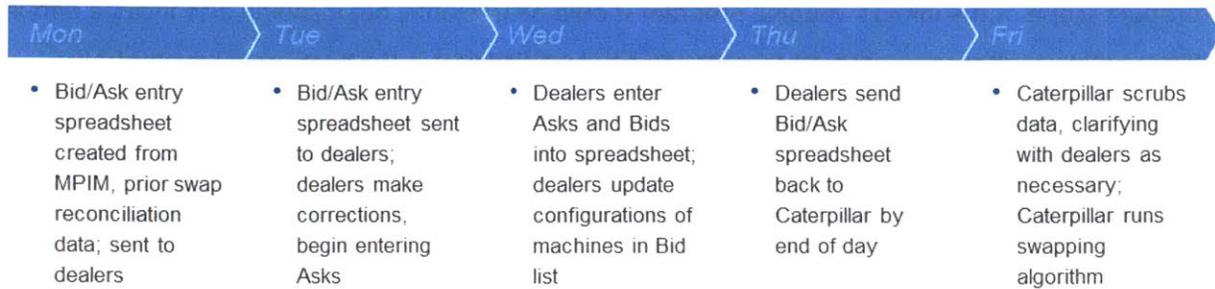
Swapping equipment from dealer to dealer creates logistical challenges with respect to the transportation, ownership, insurance, and financial reconciliation of swapped equipment. Ultimately, each of these challenges needs a robust – likely centralized – solution. However, the pilot should leverage the fact that dealers have been solving each of these problems when conducting ad hoc swaps on their own. When the pilot district has been selected, the extant process for conducting ad hoc swaps should be documented into a temporary set of standard work that can be used for the duration of the pilot.

#### 6.1.3. Pilot Process

The process used for the pilot program will be largely manual and will need to be stewarded by a dedicated resource from Caterpillar. The justification for this is that manual processes – while frequently less efficient than partially or fully automated ones – are more robust and can be adapted as problems arise. Figure 50 below presents a preliminary bi-weekly process which can be used to drive the pilot program.

## Preliminary pilot process

### Week A



### Week B



Based upon the feedback received from dealers, it's likely that the optimal cadence for a durable swapping system is weekly. The pilot cadence extends this to a two-week window that separates the processes of posting swap bids and asks and the process of initiating swap reconciliation. The justification for this dilated process is two-fold: (1) the manual pilot process will be substantially more labor intensive, necessitating more time for each cycle and (2) the longer process will provide time for problem solving as challenges arise with the pilot.

For the pilot program, information collection and transfer will likely be executed through Excel spreadsheets. Although this platform is labor intensive and may be subject to data input errors, it is accessible, has a shallow learning curve, and facilitates data input design iteration.

#### 6.1.4. Pilot Organizational Ownership

Within Caterpillar, there are two potential natural owners for the proposed swapping mechanism. The Global Supply Network Division (GSND) has historically overseen value chain efforts that cut across various product divisions. The other potential owner for a pilot program is the dealer facing commercial organization, Global Construction Industries (GCI). For the duration of the pilot, it's likely that a dedicated individual from either organization can assume the role of process manager, but someone from GCI would be preferential. Dealers typically have relatively strong relationships with

representatives in the GCI group, and this would provide a familiar touchpoint to introduce and shepherd an unfamiliar program.

## 6.2. Durable System Development

A constrained pilot swapping program will provide significant learning opportunities and will likely shape the formation of a durable swapping system. However, there are already several system development challenges which are apparent *ex ante*. These challenges and potential solutions are outlined in this section, but this framework should be adapted as information is collected in the pilot.

### 6.2.1. Inventory Monitoring System

The current inventory monitoring system, the Machine Products Inventory Manager (MPIM), represents a step change in inventory visibility over its predecessor, and the dealer-to-dealer visibility afforded by this system is a critical facilitator for the current *ad hoc* swapping that occurs within Caterpillar's dealer network. However, some of the usability and data integrity characteristics of MPIM suggest that, in its current state, MPIM will be insufficient to act as the data backbone for a durable swapping system.

From a usability standpoint, MPIM is hindered by the fact that machine configuration information is only tracked according to the assembly numbers that compose a machine. Unfortunately, these numbers can change such that multiple numbers may imply the same feature (e.g., 123-4567 and 123-4568 may both indicate that a machine has third-valve functionality), making swap specification more challenging. From a data integrity perspective, MPIM fails to track situations when dealers modify a machine after taking receipt from the factory, and anecdotal evidence suggests that some data in MPIM does not update quickly enough to support a weekly swap cadence.

### 6.2.2. Bid/Ask Posting System

For the pilot system, bids and asks will be communicated to the process manager via Excel spreadsheets. However, the durable system will require some type of web entry system that streamlines data entry, helps to control user data input errors, and communicates data directly to the matching algorithm. The exact format of the bid / ask posting system will be informed by the pilot, but it should be incorporated as a part of Caterpillar's next generation dealer equipment ordering system such that dealers can consider multiple procurement routes (e.g., factory purchase vs. dealer swap) in parallel.

### 6.2.3. Matching System

The pilot matching system is currently composed of four separate Python scripts which are oriented towards simulation as opposed to practical match execution. While the core matching portion of the script can be maintained, it can be substantially streamlined. Furthermore, the durable matching system should be integrated with critical data sources (i.e., MPIM's next generation and the bid/ask system) and

a graphical user interface should be developed to improve usability. Caterpillar's internal analytics functions are well-equipped to support this system development.

#### 6.2.4. Freight System

In order to ensure that swaps are physically reconciled in a timely fashion, the freight system used to move machines from one dealer to another in the swapping program should also be managed centrally; Caterpillar's Last Mile network is the logical owner for this freight system. Although volume of machine shipments required by the proposed swapping program is much smaller than the factory to dealer shipments currently managed by Caterpillar's Last Mile program, it still presents novel challenges.

First, current Last Mile shipments in North America typically originate from one of only several plant locations (or ports) across the United States. In the proposed swapping program, a given swap chain could originate at any of Caterpillar's dealers within the ADSD-N region. Furthermore, while Caterpillar's current Last Mile shipments are largely predictable, both from a timing and destination standpoint, the shipments required by a swapping program will vary week to week both with respect to the origin-destination pairs and the volume. Both of these idiosyncrasies complicate the scheduling of machine transport. Conversely, the swapping program presents an opportunity for improved backhaul. As long as the machines being swapped across a given chain are of relatively similar masses and dimensions, a single truck may be able to achieve full utilization across a given chain, dropping off and picking up a machine at each destination.

During the pilot, a Last Mile representative should be invited to observe the system and collect data that will be needed to develop the processes required to support a durable swapping logistics system.

#### 6.2.5. Financial System

One of the challenges posed by a swapping system in which a dealer may be interfacing with more than one counterpart in a single swap (receiving a machine from one dealer, sending a machine to a separate dealer) is the financial reconciliation that must occur when those machines do not have the same wholesale value. Under the current ad-hoc system, each of these transactions would have to be determined and reconciled separately. Once again, this process could be streamlined with a centralized system; the system would track the wholesale value of each machine being swapped and would bill or reimburse each dealer according to the difference in value of the machine that was received in a swap versus the machine that was sent.

The challenges for developing such a centralized system are at least two-fold. First, the system must be able to communicate with each dealer's financial system. This could present an intractable barrier due to the fact that each dealer has implemented its own financial systems. Second, the proposed central

financial clearing system would have to have accurate wholesale cost information on each machine being swapped; this could be particularly challenging in cases where a dealer has changed options on a machine, even if those changes are as simple as swapping tire sets.

#### 6.2.6. Legal Framework

Finally, a durable, centralized swapping program would need some legal framework to operate within. Questions such as when each dealer (or Caterpillar itself) relinquishes and receives ownership of each machine involved in a swap and who is responsible for insuring the machine must be answered. It's possible that the swapping program would need to be executed through a new entity within the Caterpillar ecosystem.

GSND is likely the best starting point for resolving the challenge of developing an appropriate legal and financial framework for a durable, centralized swapping program.

#### 6.3. Durable Process Development

In addition to the systems that will underpin a durable swapping network, the process discussed in Section 6.1.3 must be refined and formalized. The steady-state process will be heavily contingent upon the results of the pilot and the capabilities of the systems that are developed. Nevertheless, the durable process will likely have the following characteristics:

- The process will **proceed on a weekly cadence**. Preliminary conversations with dealers suggest that a cycle in which dealers finalize bids/asks on Friday and receive swap matches on Monday would minimize the risk of disruption to their current sales processes.
- At least initially, the process must **include a period in which dealers can review and veto swaps**. Although a veto could disrupt a potentially valuable swap, the presence of a 'safety valve' for dealers will help increase participation and comfort with the system.
- Even if the durable systems afford a largely automated process, the initial steady-state process should **include human touch by the GCI group** (e.g., GCI should review the automatically identified swaps before sending them to dealers). This will help prevent obviously unviable swaps that could arise from data or system errors, and it will also help build dealer and Caterpillar confidence in the system.

Each of the above considerations are designed to make the swapping process attractive to dealers. Ultimately, the most critical consideration in architecting the periodic activities associated with a swapping program is ensuring that all of the participant dealers trust the process.

## 7. Conclusions and Extensions

Overall, this thesis suggests that network swapping is a potentially compelling technique to capture some of the benefits of inventory pooling in circumstances where fragmented ownership would otherwise prevent virtual pooling. The approach presented in this thesis is a simple first attempt at developing a network swapping concept, and there are a number of academic extensions of this problem which could further refine the concept and improve impact.

### 7.1. Conclusions

The results observed from the network swapping simulation yield several notable conclusions.

- **The potential benefit from network swapping is significant.** Every scenario examined provided significant operational benefits in terms of inventory reduction, backorder reduction, and sales acceleration. With the exception of the two-month lead time scenario, every scenario was also NPV positive, with dealer benefits ranging approximately from USD 3 to 5 MM in likely scenarios (before accounting for any ongoing revenue benefits from incremental captured sales).
- **Benefits can be realized even without a nationwide network.** The scenario in which only intra-district swaps were allowed still showed operational and financial benefits, although operational benefits were attenuated. This result is important because it suggests that a staged rollout in which swapping networks are first implemented within marketing districts is still financially attractive.
- **The benefits of network swapping are heterogeneous across dealers.** Although almost every dealer observes positive operational and financial impacts from participating in network swapping, those benefits are heterogeneous with some dealers observing significantly higher improvements. Such asymmetry could present a risk to a swapping network – dealers observing a lower return from the swapping network may drop out – and some mechanism in which the financial benefits of network swapping can be shared should be considered. Future generations of the swapping simulation could simulate the effect of potential drop-outs.
- **Network swapping provides a greater benefit under periods of supply chain constraint.** The factory lead time sensitivity analysis demonstrates that network swapping provides greater incremental benefits when lead times dilate. This result suggests that standing up a network swapping mechanism could provide additional benefits as Caterpillar and its dealers are exposed to the cyclical industry demand that is characteristic of construction equipment.
- **Logistics costs from network swapping are significant and must be analyzed further.** The discounted cost of logistics under the proposed network swapping system consume a significant portion of the financial benefit associated with accelerated sales and reduced inventory holding

costs, and the cost of logistics also has a significant impact on the relative benefit of different scenarios. The simple optimization proposed in this thesis implicitly includes logistics cost in the constraints as opposed to the objective function; this likely is not viable for a durable swapping algorithm. Finding a way to migrate this cost to the objective function could significantly improve financial results.

- **The potential benefits of network swapping are relatively robust to the configuration specifications imposed by dealers.** The sensitivity analysis conducted on configuration complexity demonstrates that the financial and operational benefits do not decay rapidly as more restrictions in configuration are imposed. Indeed, as a consequence of secondary logistics cost impact, the financial benefit actually improves with higher specificity. This is a promising result which suggests that dealers who wish to specify the machines they receive in a swap along many dimensions will not dramatically impact the efficacy of network swapping.

## 7.2. Proposed Project Extensions

The results of this thesis suggest that the core methodology proposed for network swapping can be extended and improved. Furthermore, there are a number of ancillary problems which could be examined to augment the impact of network swapping.

### 7.2.1. Network Extension

The results presented in this thesis only consider a single machine class: Medium Wheel Loaders (MWLs). The reason that this thesis is circumscribed to this machine class is a function of scope management and does not reflect the actual applicability of a network swapping mechanism. In reality, such an initiative should *at least* extend to all machines within the general Construction Industries (CI) class. The benefit of this extension should be super-linear: since machines can be swapped *across* classes as well as within classes, the number of available swaps will scale super-linearly (e.g., if a second machine class with the same sales volumes as MWLs were added, the observed benefits should more than double).

Furthermore, it's likely that the proposed swapping mechanism could be extended to include smaller Building and Construction Product (BCP) machines (e.g., Compact Wheel Loaders), provided that such machines could be traded in a 'bundle' for the larger CI machines to maintain approximate value parity in the swap. This could be done under the current swap identification algorithm by defining the requisite bundles ex-ante, or an alternative swap identification algorithm could be designed which allows dealers to specify their own target bundles.

### 7.2.2. Alternative Swap Identification Techniques

The methodology used to identify potential swaps in this thesis is a simple longest-path optimization. There are several problems associated with this optimization methodology: the solved objective function is merely a proxy for the true underlying objective function (i.e., net present value improvement), the optimization does not allow dealers to express any preference among several potentially viable swaps, and the optimization has several hard constraints which may be loosened in real world application (e.g., the conservation of mass constraint which states that a dealer must receive a machine in order to send a machine to another dealer).

There are several options to improve this swap identification methodology. If heuristics can be designed to approximate the value of each individual swap edge, an objective function could be designed which solves for maximum incremental value as opposed to maximum inventory movement. Second, there are heuristics such as the random serialized dictatorship which could allow dealers to express a preference between several potentially viable swaps. Those same heuristics could also be designed in a way which allows for certain hard constraints, such as the conservation of mass constraint, to be loosened (in the random serialized dictatorship, for example, a requesting dealer could preference 'no machine' above a portion of its own inventory; if another dealer sought a unit in that deprioritized block, it could receive the machine without sending one of its own).

### 7.2.3. Complementary Initiative Impact

Caterpillar has several complementary Engineered Value Chain (EVC) initiatives whose potential impact on network swapping has not been fully explored. Chief among these are delayed differentiation initiatives such as Finish to Order (FTO) and dealer kitting. These delayed differentiation initiatives add direct cost to each machine but also function to lessen potential constraints around dealer configuration complexity requirements. Dealers currently have varying strategies around how to manage kitting and delayed differentiation (some dealers inventory kits and use them frequently while others do not use them at all); an analysis of these initiatives in conjunction with a network swapping simulation could provide guidance on how best to pursue a combined implementation.

### 7.2.4. Multi-echelon Problem

The presence of Caterpillar's Product Distribution Centers (PDCs) implies that the network swapping problem could be framed within a larger multi-echelon optimization problem. Addressing such a problem could suggest alternative stocking levels at the dealers and PDCs, particularly given the opportunity to differentiate machines at the PDC using FTO. This problem once again raises issues of agency: the step from the PDC echelon to the dealer echelon crosses an ownership barrier. However,

this barrier may not be insurmountable if a significant opportunity can be identified by rebalancing relative inventory levels.

#### 7.2.5. Logistics Network Problem

In this thesis, the logistics cost of executing a swap for a specific machine along a specific edge were taken as a constant (based upon quotes from Caterpillar's internal Last Mile group), and it was implicitly assumed that all of the identified swaps would take one month to execute. In reality, the logistics associated with executing a network swapping mechanism are their own optimization problem with the potential to significantly affect the financial impact and feasibility of network swapping. For example, if a single truck can be used to execute a chained series of swaps, network swapping has the potential to significantly increase backhaul and reduce per-mile costs of transporting equipment. However, depending upon the machines that are being moved along each edge, a single truck may not be able to execute every swap. Furthermore, for longer swap chains, it may not be feasible to execute every swap within the prescribed time period unless multiple trucks are used.

## References

- [1] R. Chernow, *Titan : the life of John D. Rockefeller, Sr.* .
- [2] B. Tillman, "D-Day Logistics: Preparing for Landfall - History," *History on the Net*. [Online]. Available: <https://www.historyonthenet.com/d-day-logistics/>. [Accessed: 11-Jan-2018].
- [3] M. Bellis, "History of the Bulldozer - Famous Inventions," *ThoughtCo*, 2017. [Online]. Available: <https://www.thoughtco.com/history-of-the-bulldozer-1991353>. [Accessed: 12-Jan-2018].
- [4] "Caterpillar | 2016 Annual Report to Shareholders on Form 10-K," *SEC Form 10-K*, 2016. [Online]. Available: <https://www.caterpillar.com/en/investors/financialInfo/proxy-statement/general-and-financial-information-proxy-appendix.html#s4687FA22158CF121ED52D894A9A7B41E>. [Accessed: 12-Jan-2018].
- [5] M. G. Doyle, "KHL: Global construction equipment market down 16.2%; Cat retains no. 1 spot," *Equipment World*, 2016. [Online]. Available: <https://www.equipmentworld.com/khl-global-construction-equipment-market-down-16-2-cat-retains-no-1-spot/>. [Accessed: 11-Jan-2018].
- [6] D. V Fites, "Make your dealers your partners," *Harv. Bus. Rev.*, vol. 74, no. April, p. 84, 1996.
- [7] J. B. Kelleher, "CATERPILLAR: Our Dealers Are Missing Up To \$18 Billion In Easy Sales - Business Insider," *Reuters*, 2014. [Online]. Available: <http://www.businessinsider.com/r-insight-from-dumb-iron-to-big-data-caterpillars-dealer-sales-push-2014-20>. [Accessed: 12-Jan-2018].
- [8] G. P. Cachon and C. Terwiesch, *Matching Supply with Demand: An Introduction to Operations Management*. New York: McGraw-Hill, 2008.
- [9] "Transplant trends | UNOS." [Online]. Available: [https://unos.org/data/transplant-trends/#transplants\\_by\\_organ\\_type+year+2017](https://unos.org/data/transplant-trends/#transplants_by_organ_type+year+2017). [Accessed: 15-Apr-2018].
- [10] P. Starzl, Thomas E., MD, "Felix T. Rapaport, MD," *J. Am. Med. Assoc.*, vol. 285, no. 23, p. 3032, 2001.
- [11] B. C. Wallis, K. P. Samy, A. E. Roth, and M. A. Rees, "Kidney paired donation 2011.," *Nephrol Dial Transpl.*, vol. 26, pp. 2091–2099, 2011.
- [12] D. Serur and M. Charlton, "Kidney paired donation 2011.," *Prog. Transplant.*, vol. 21, no. 3, pp. 215–218, 2011.
- [13] K. Park, J. Il Moon, S. Il Kim, and Y. S. Kim, "Exchange donor program in kidney transplantation," *Transplantation*, vol. 67, no. 2, pp. 336–338, 1999.
- [14] B. Ellison, "A Systematic Review of Kidney Paired Donation: Applying Lessons From Historic and Contemporary Case Studies to Improve the US Model," *Whart. Res. Sch.*, vol. 107, 2014.
- [15] "First kidney paired donor transplants performed as part of national pilot program," *United Network for Organ Sharing*, 2010. .
- [16] A. E. Roth, T. Sönmez, and M. U. Ünver, "A kidney exchange clearinghouse in New England," *Am. Econ. Rev.*, vol. 95, no. 2, pp. 376–380, 2005.
- [17] K. Sack, "60 Lives , 30 Kidneys , All Linked," *New York Times*, New York City, pp. 1–18, 18-Feb-

2012.

- [18] M. L. Melcher *et al.*, “Kidney transplant chains amplify benefit of nondirected donors,” *JAMA Surg.*, vol. 148, no. 2, pp. 165–169, 2013.
- [19] G. D. Eppen, “Effects of centralization on expected costs in a multi-location newsboy problem,” *Manage. Sci.*, vol. 25, no. 5, pp. 498–501, 1979.
- [20] D. Smirnov and Y. Gerchak, “Inventory sharing via circular bidirectional chaining,” *Int. J. Prod. Econ.*, vol. 179, no. 2, pp. 141–152, 2016.
- [21] C. Paterson, G. Kiesmüller, R. Teunter, and K. Glazebrook, “Inventory models with lateral transshipments: A review,” *Eur. J. Oper. Res.*, vol. 210, no. 2, pp. 125–136, 2011.
- [22] K. Bimpikis and M. G. Markakis, “Inventory Pooling Under Heavy-Tailed Demand,” *Manage. Sci.*, vol. 62, no. 6, pp. 1800–1813, 2016.
- [23] T. C. K. Mark L. Berenson, David M. Levine, *Basic Business Statistics: Concepts and Applications*, 12th ed. Boston: Prentice Hall, 2012.
- [24] J. J. Neale and S. P. Willems, “The failure of practical intuition: How forward-coverage inventory targets cause the landslide effect,” *Prod. Oper. Manag.*, vol. 24, no. 4, pp. 535–546, 2015.
- [25] L. Taccari, “Integer programming formulations for the elementary shortest path problem,” *Eur. J. Oper. Res.*, 2014.

## Appendix

### Executive Review Slides



**Virtual Pooling Approximation Using  
Longest Path Network Optimization**

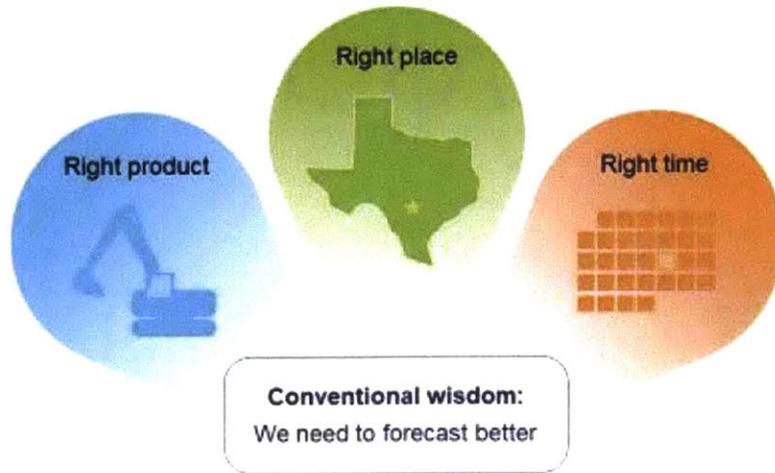
**What defines a good supply chain?**

*We must engineer our value chain, much like we would engineer a product. **Starting with the customer requirements** and designing the entire value chain based on those requirements all the way back to the engineer's pen.*

SOURCE: Caterpillar Intranet. *Engineered Value Chain*



## What defines a good supply chain?



Caterpillar plans using an aggregate demand profile; this demand has variability but is forecastable...



SOURCE: Caterpillar BIC, SSIO 10-year data

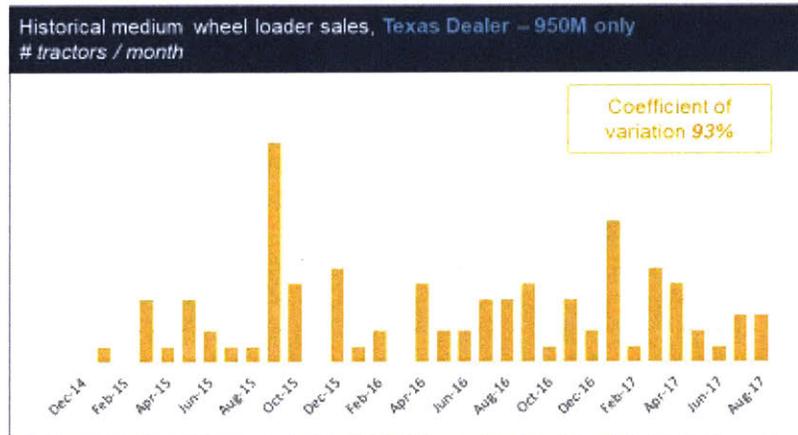
... however, this ability to reasonably forecast demand *does not* translate down to the point of sale



SOURCE: Caterpillar BIC, SSIO 10-year data

5

... however, this ability to reasonably forecast demand *does not* translate down to the point of sale



SOURCE: Caterpillar BIC, SSIO 10-year data

6

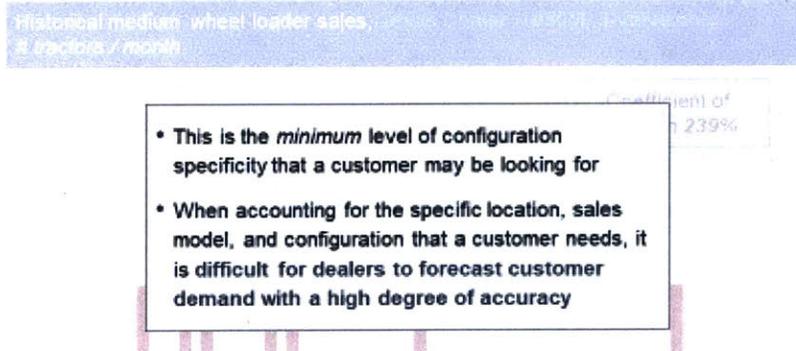
... however, this ability to reasonably forecast demand *does not* translate down to the point of sale



SOURCE: Caterpillar BIC, SSIO 10-year data

7

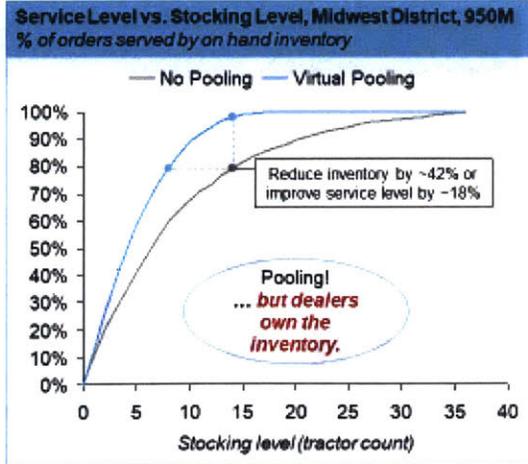
... however, this ability to reasonably forecast demand *does not* translate down to the point of sale



SOURCE: Caterpillar BIC, SSIO 10-year data

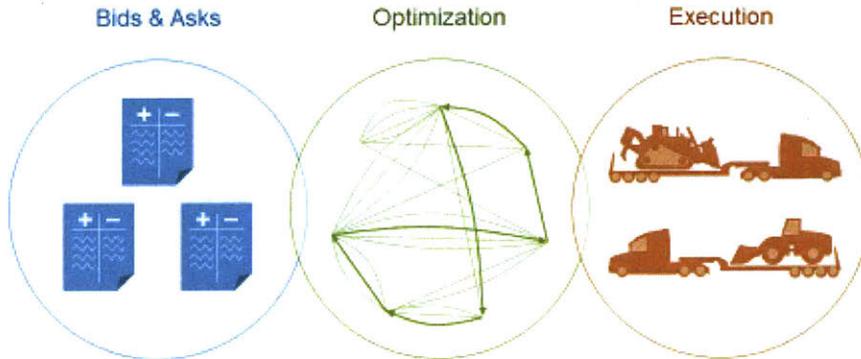
8

CAT can counter this challenge through virtual pooling, but its commercial strategy is a barrier



SOURCE: Caterpillar BIC, SSIO 10-year data

This barrier can be overcome with a centrally managed dealer-to-dealer swapping network



Consider the following scenario in which CemCo is seeking to purchase two new wheel loaders



**TEXAS CAT**

- CemCo has decided to expand a quarry in West Texas and needs two additional 950 sized wheel loaders
- CemCo's application requires 3-valve functionality
- While Texas CAT typically stocks at least six 950s, it sells very few 3-valve machines and only has one in inventory
- Current factory availability is 14 weeks
- CemCo is under a tight time constraint and has had success with Volvo loaders in other regions

11

Texas CAT is able to source the 3-valve equipped 950M it needs by sending a 966M to North Carolina

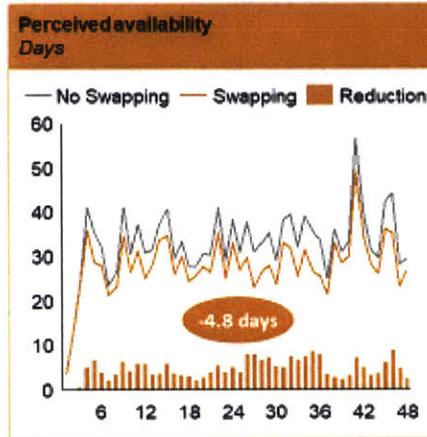
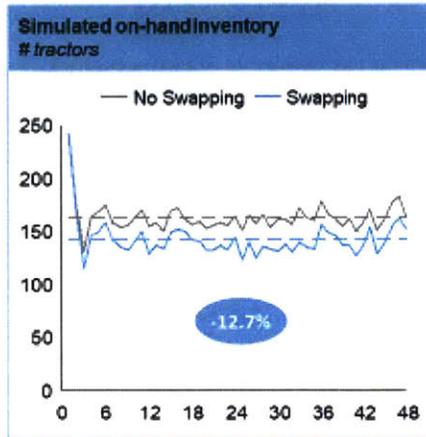


- Texas CAT participates in a four edge swap with Massachusetts, New Jersey, and North Carolina
- Texas CAT receives a standard lift, 3-valve, manual diff 950M from New Jersey to make the sale to CemCo
- Texas CAT sends a standard lift, 2-valve, manual diff 966M for Carolina to make a sale against one of its own backorders

12

This enhanced swapping<sup>1</sup> decreases on-hand inventory and improves perceived availability

EASTERN NETWORK, MWL ONLY



<sup>1</sup> Simulation for 27 dealers + PDC. includes all M-series MWL models. critical swap parameters are hydraulics, differential, and boom

SOURCE: CAT Swapping Model v2.0, 11.14.2017 model run

13

Additional research could be pursued to further the network swapping concept



**Alternative algorithms**

- Hard 'conservation of mass' constraint may be too strict
- Random serialized dictatorship is promising heuristic



**Complementary initiatives**

- Impact of delayed differentiation on swap density
- Optimal PDC stocking level under swapping condition



**'Non-iron' swapping**

- Integration of build slots into swapping algorithm
- Build slot 'options' trading

14

## EVC Initiative Idea List

 <p><b>Push to FTO</b></p>	<ul style="list-style-type: none"> <li>• What additional product options can be pushed down to the third-party PDC (third-valve, axle oil cooler, auto-lube, etc.)?</li> <li>• In addition to big questions above, how does this affect quality and delivery time?</li> </ul>
 <p><b>Push to kitting</b></p>	<ul style="list-style-type: none"> <li>• What product options can be offered strictly as kits dealer install?</li> <li>• Can CAT offer standard machines which can be made into full offer range at the dealer?</li> <li>• Would dealers be willing to accept this if it resulted in shorter lead times / more control?</li> </ul>
 <p><b>Delayed differentiation</b></p>	<ul style="list-style-type: none"> <li>• How early is the complexity of certain options introduced into the manufacturing process? Can this complexity be introduced later?</li> <li>• Common frame analysis would fall under this category</li> </ul>
 <p><b>Inventory pooling</b></p>	<ul style="list-style-type: none"> <li>• Can we aggregate inventory at a higher level to alleviate individual dealer stochasticity?</li> <li>• Can we more efficiently create virtual pooling across dealers and PDCs without physically pooling inventory (e.g. longest chain matching for inventory trades &amp; build slots)?</li> </ul>
 <p><b>Best pathway delivery</b></p>	<ul style="list-style-type: none"> <li>• When dealers order a product, can we increase visibility with respect to all possible delivery pathways (e.g. pull from dealer A with modifications 1 &amp; 2, pull from PDC with modification 3, or order from CAT)? Allow dealers to make cost/lead time trade off</li> </ul>
 <p><b>Bundling / option rationalization</b></p>	<ul style="list-style-type: none"> <li>• Can any options be more efficiently bundled or offered as standard to reduce complexity?</li> <li>• Can any options be rationalized?</li> <li>• Is there an algorithm/process that can be established for determining cost of commonality?</li> </ul>
 <p><b>Design for complexity</b></p>	<ul style="list-style-type: none"> <li>• Can the machine be redesigned to more easily accept the complexity of the product?</li> <li>• Could a redesign form a foundation for additional FTO / kitting in the future?</li> </ul>
 <p><b>PDC to dealer inventory shift</b></p>	<ul style="list-style-type: none"> <li>• Can inventory be shifted from PDCs to dealers to reduce holding cost for CAT and increase machine accessibility for the dealers?</li> <li>• How does this affect our ability to pool inventory (more fragmented, but virtually?)?</li> </ul>

 <p><b>Build slot allocation optimization</b></p>	<ul style="list-style-type: none"> <li>• During periods of severe constraint (i.e. high demand and low production capacity), <b>how should CAT allocate build slots to its dealers?</b></li> <li>• Can build slots be 'valued' by dealers to determine correct rank?</li> </ul>
 <p><b>Crate delivery</b></p>	<ul style="list-style-type: none"> <li>• Can more <b>machines be delivered to PDC or to the dealers in partially assembled form</b> (i.e. in shipping crates) to reduce freight and facilitate inventory sharing?</li> <li>• How would this work in conjunction with kitting / FTO?</li> </ul>
 <p><b>Feature driven price list</b></p>	<ul style="list-style-type: none"> <li>• The current price list is largely a product of the way the BOM for a machine is structured; dealers have to worry about ordering conflicting components on a machine</li> <li>• <b>Make ordering feature driven</b> (more like automotive)?</li> </ul>
 <p><b>Alternative sourcing optimization</b></p>	<ul style="list-style-type: none"> <li>• A tradeoff exists between the lower machine costs that can be realized from alternative sourcing (e.g. manufacturing in China) and the complexity introduced into S&amp;OP</li> <li>• What is the <b>optimal amount of alternative sourcing into each region?</b></li> </ul>
 <p><b>Line-side cab finishing</b></p>	<ul style="list-style-type: none"> <li>• Currently, cabs are the most complex sub-system within medium wheel loaders; completed cabs are delivered in sequence from a separate CAT facility</li> <li>• Can CAT <b>receive simpler, unfinished cabs and finish them to spec line-side?</b></li> </ul>
 <p><b>Parts inventory pooling</b></p>	<ul style="list-style-type: none"> <li>• For machines in production, there will be WIP inventory of parts at the plant and replacement parts inventory in Morton</li> <li>• Should this <b>parts inventory all be pooled at the manufacturing facility itself?</b></li> </ul>
 <p><b>Cost of complexity (\$)</b></p>	<ul style="list-style-type: none"> <li>• What is the <b>full OPACC impact of introducing complexity</b> in the line? In the PDC? At the dealer? Is there a process we can follow to determine this cost for future options?</li> <li>• This would probably need to be scoped down to a single system or option set</li> </ul>
 <p><b>Value of availability +\$</b></p>	<ul style="list-style-type: none"> <li>• What is the <b>value of being able to deliver machines to dealers and customers within a given time frame?</b> What is the value of being able to precisely forecast the delivery time?</li> <li>• Perhaps this decision should be dealer driven</li> </ul>



Cost of  
availability

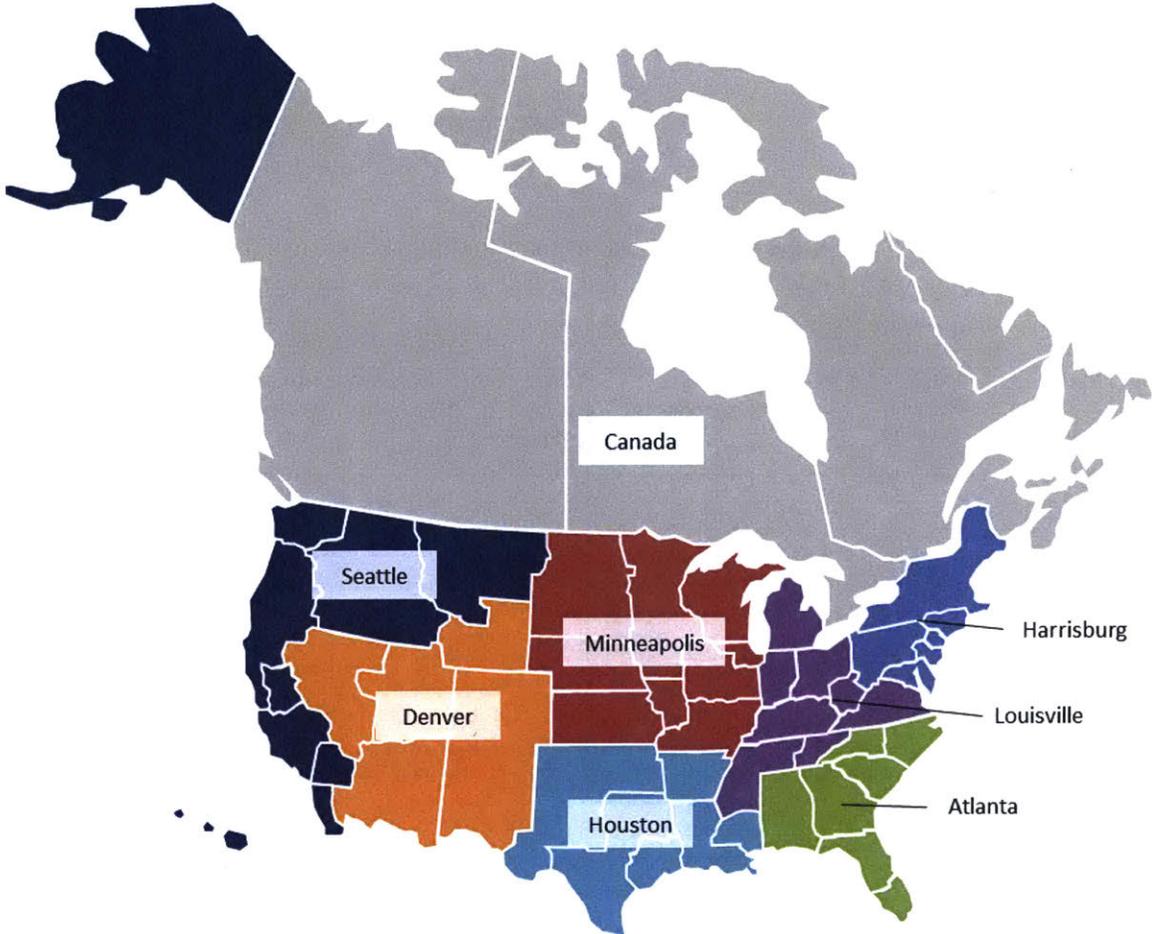
- What is the **cost of delivering a machine in a given time frame**? How much does it cost to accelerate the delivery of a machine by one week; how much can be saved by delaying the delivery of a machine by one week?

### Service Level Sample Analysis

The table below shows the period by period results of the simplified base stock model used to approximate the dealer to customer service level in Caterpillar's dealer network. The results presented below are for a single dealer / sales model point of sale.

Period	Factory Deliveries	Period Beginning Inventory	New Customer Orders	Sales Against Backorders	Sales Against New Orders	Backorder Additions	Period Ending Backorders	Period Ending Inventory	Factory Orders
1	0	7	0	0	0	0	0	7	0
2	0	7	1	0	1	0	0	6	1
3	0	6	1	0	1	0	0	5	1
4	0	5	2	0	2	0	0	3	2
5	1	4	3	0	3	0	0	1	3
6	1	2	4	0	2	2	2	0	4
7	2	2	3	2	0	3	3	0	3
8	3	3	1	3	0	1	1	0	1
9	4	4	1	1	1	0	0	2	1
10	3	5	5	0	5	0	0	0	5
11	1	1	2	0	1	1	1	0	2
12	1	1	0	1	0	0	0	0	0
13	5	5	2	0	2	0	0	3	2
14	2	5	2	0	2	0	0	3	2
15	0	3	0	0	0	0	0	3	0
16	2	5	0	0	0	0	0	5	0
17	2	7	0	0	0	0	0	7	0
18	0	7	2	0	2	0	0	5	2
19	0	5	0	0	0	0	0	5	0
20	0	5	4	0	4	0	0	1	4
21	2	3	3	0	3	0	0	0	3
22	0	0	0	0	0	0	0	0	0
23	4	4	0	0	0	0	0	4	0
24	3	7	1	0	1	0	0	6	1
25	0	6	0	0	0	0	0	6	0
26	0	6	0	0	0	0	0	6	0
27	1	7	0	0	0	0	0	7	0
28	0	7	0	0	0	0	0	7	0
29	0	7	4	0	4	0	0	3	4
30	0	3	1	0	1	0	0	2	1
31	0	2	0	0	0	0	0	2	0

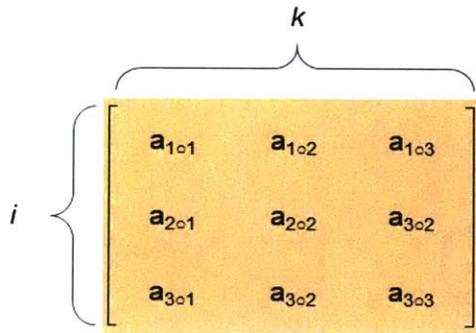
Dealer District Map



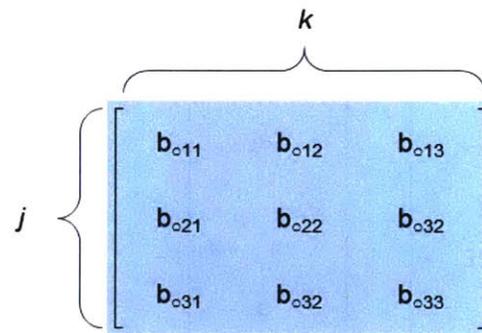
### Row-Wise Product Description

The 'row-wise' product discussed in section 5.2.4 is an element wise product of two orthogonal matrices which share a single dimension. Consider the  $ik$  and  $jk$  dimensioned matrices below, whose elements have been written in  $ijk$  space. The row-wise product is taken such that the  $ijk$ 'th element of the product matrix is the element product of the  $ik$ 'th element of the first matrix and the  $jk$ 'th element of the second matrix. Consequently, the resulting matrix is three-dimensional.

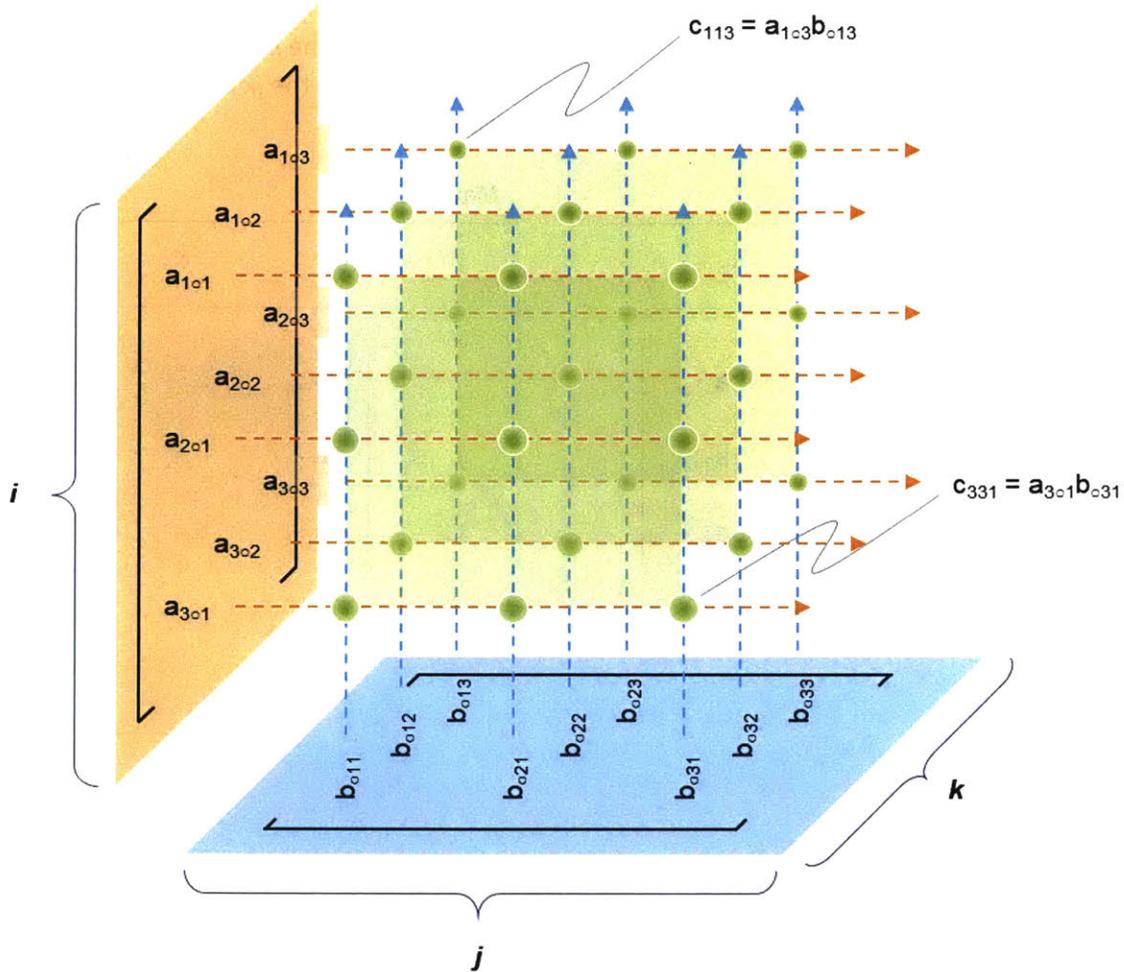
**Matrix A**



**Matrix B**



## Row-wise product



## Matrix C

$c_{111} = a_{1o1}b_{o11}$	$c_{121} = a_{1o1}b_{o21}$	$c_{131} = a_{1o1}b_{o31}$	$c_{112} = a_{1o2}b_{o12}$	$c_{122} = a_{1o2}b_{o22}$	$c_{132} = a_{1o2}b_{o32}$	$c_{113} = a_{1o3}b_{o13}$	$c_{123} = a_{1o3}b_{o23}$	$c_{133} = a_{1o3}b_{o33}$
$c_{211} = a_{2o1}b_{o11}$	$c_{221} = a_{2o1}b_{o21}$	$c_{231} = a_{2o1}b_{o31}$	$c_{212} = a_{2o2}b_{o12}$	$c_{222} = a_{2o2}b_{o22}$	$c_{232} = a_{2o2}b_{o32}$	$c_{213} = a_{2o3}b_{o13}$	$c_{223} = a_{2o3}b_{o23}$	$c_{233} = a_{2o3}b_{o33}$
$c_{311} = a_{3o1}b_{o11}$	$c_{321} = a_{3o1}b_{o21}$	$c_{331} = a_{3o1}b_{o31}$	$c_{312} = a_{3o2}b_{o12}$	$c_{322} = a_{3o2}b_{o22}$	$c_{332} = a_{3o2}b_{o32}$	$c_{313} = a_{3o3}b_{o13}$	$c_{323} = a_{3o3}b_{o23}$	$c_{333} = a_{3o3}b_{o33}$

## Baseline Demand Scenarios, Gross Description

The following tables characterize the seven demand runs which were used to generate the Baseline model results. The figures in the tables below are expressed as percentages of the expected mean sales in each category to disguise Caterpillar's demand (i.e., if a given model had a mean expected demand of

100 machines over the 48-month period and run 1 produced a demand of 105 machines, the figure in the corresponding square would be 105%).

Simulated demand characteristics by Caterpillar marketing district:

	Run 1	Run 2	Run 3	Run 4	Run 5	Run 6	Run 7	Average
Harrisburg	93.8%	99.8%	98.2%	95.8%	97.6%	98.0%	99.3%	97.5%
Louisville	104.8%	97.0%	100.4%	97.4%	105.4%	90.5%	103.8%	99.9%
Atlanta	101.4%	104.4%	102.6%	98.8%	100.6%	102.7%	101.1%	101.7%
Minneapolis	104.5%	99.2%	98.9%	104.7%	96.7%	104.4%	97.8%	100.9%
Denver	101.3%	95.6%	104.7%	106.3%	98.6%	95.2%	96.6%	99.8%
Seattle	104.2%	94.5%	91.7%	101.5%	94.5%	99.4%	102.6%	98.4%
Canada	103.9%	104.0%	100.2%	105.5%	95.1%	99.3%	107.2%	102.2%
<b>Total</b>	<b>100.8%</b>	<b>99.1%</b>	<b>99.0%</b>	<b>100.8%</b>	<b>97.9%</b>	<b>98.7%</b>	<b>100.8%</b>	<b>99.6%</b>

Simulated demand characteristics by sales model:

	Run 1	Run 2	Run 3	Run 4	Run 5	Run 6	Run 7	Average
950M	99.2%	100.1%	98.6%	101.4%	96.2%	96.4%	100.2%	98.9%
962M	97.7%	106.5%	95.2%	101.7%	99.1%	106.0%	103.7%	101.4%
966M	98.1%	105.1%	101.0%	101.2%	100.1%	98.0%	101.5%	100.7%
966MXE	110.0%	95.0%	82.5%	107.5%	101.7%	95.8%	105.0%	99.6%
972M	101.1%	96.5%	96.7%	101.6%	95.1%	97.4%	99.0%	98.2%
972MXE	103.6%	88.5%	108.4%	111.2%	76.9%	96.8%	93.3%	97.0%
980M	102.8%	95.7%	101.6%	102.0%	97.5%	101.1%	104.6%	100.8%
982M	106.1%	94.6%	95.7%	89.8%	107.4%	101.7%	93.1%	98.4%
<b>Total</b>	<b>100.8%</b>	<b>99.1%</b>	<b>99.0%</b>	<b>100.8%</b>	<b>97.9%</b>	<b>98.7%</b>	<b>100.8%</b>	<b>99.6%</b>