Creating Transportation Policy in a Network that utilizes both Contract Carriers and an Internally Managed Fleet

by

Michael Jay Mulqueen, Jr.

Bachelor of Science in Marketing
University of Maryland

Submitted to the Engineering Systems Division in Partial Fulfillment of the Requirements for the Degree of

MASTER OF ENGINEERING IN LOGISTICS

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

JUNE 2006

Signature of Author ....
Engineering Systems Division
May, 2006

Certified by ....
Christopher G. Caplice
Executive Director – Master of Engineering in Logistics
Thesis Supervisor

Accepted by ....
Yosef Sheffi
Professor, Engineering Systems Division
Professor, Civil and Environmental Engineering Department
Director, MIT Center for Transportation and Logistics

Creating Transportation Policy in a Network that utilizes both Contract Carriers and an Internally Managed Fleet

by

Michael Jay Mulqueen, Jr.

Submitted to the Engineering Systems Division in Partial Fulfillment of the Requirements for the Degree of
MASTER OF ENGINEERING IN LOGISTICS

ABSTRACT

A convergence of factors including a strong economy, changing demographics and increased regulatory control has resulted in a U.S. For-Hire Truckload (TL) industry that is increasingly pressed to meet shippers' needs for freight services. As demand now exceeds supply, the buyer/seller relationship has swung to favor TL providers who wield much more power than they have historically enjoyed. TL carriers are now able to negotiate higher rates, increase charges for providing additional services and decrease service levels as shippers are unable to find suitable replacement carriers.

Many shippers have responded to these changing market dynamics by increasing the use of private and/or dedicated fleets within their distribution networks. This provides them with guaranteed capacity, increased leverage with carriers during rate negotiations and increased overall operational control of their networks.

In this thesis, I will propose a methodology for the creation of a shipper's overall transportation policy in a distribution network that uses an internally managed fleet in conjunction with TL contract carriers

This approach constructs transportation policy in a manner that recognizes the differences in costing between an internally owned and managed fleet versus that of a contract TL carrier. It seeks to maximize savings by leveraging internal economies of scope through the assignment of fleet resources to closed loop tours. Additionally, the approach will go beyond the standard deterministic methods that are commonly employed in the creation of transportation policy. Instead, an iterative process that incorporates both optimization and simulation is proposed that ensures variability inherent within the network is taken into account when defining the best transportation policy for an organization.

Introduction

Many large shippers utilize a combination of privately owned and internally managed fleet resources in conjunction with common carriers to execute the distribution of raw materials, work in process and finished goods throughout their networks. In a recent survey, 41% of organizations stated they use a blended approach whereby the private fleet freight capacity is supplemented by the use of dedicated fleets, for-hire carriers and/or owner operators (Terreri, 2006). These enterprises translate their anticipated shipping requirements into a transportation policy, which is then used as the basis for daily execution. However, the tools and methodologies commonly used to create the overarching transportation policy do not adequately address key decision criteria that organizations which rely on a private fleet and contract carriage should be considering.

In this thesis, I will define a methodological approach that enhances the existing process of defining transportation policy for organizations that utilize a blended approach to transportation. The extensions I am proposing are:

1. The use of non-deterministic means to account for variability in a distribution network.
2. Accounting for the fixed cost and perishable nature of fleet assets
3. Introduction to the use of spot markets to increase productivity by selling excess capacity on the open market

Defining Transportation Policy Decisions

When creating transportation policy for a complex network that utilizes both for-hire and private fleets, enterprises struggle with answering the following questions:

- How big should my private fleet be?
- What lanes should my fleet cover vs. being outsourced to common carriers?
- Where do I domicile drivers?
- How do I manage demand variability and seasonality

To answer these questions, we need to understand key points that drive an organization's decisions regarding when and where to use an internally manage private fleet versus when to use For-Hire transportation. In many industries such as retail, grocery retail and wholesale distribution, there is a strong belief that only a private or dedicated fleet will provide the level of service and control necessary for customer/store deliveries. In these industries, the majority of outbound freight is executed by an internally managed fleet. These industries have addressed the above decision points by adhering to a general rule that states that inbound freight is to be handled by a For-Hire carrier, often times arranged and paid for by the manufacturer. Outbound freight is to be handled by an internal fleet. While it is unlikely that store deliveries will utilize For-Hire carriers in the future, the movement of wholesale distributors and retailers managing their inbound freight continues to gain momentum.

In terms of manufacturing, while there is a bias towards the use of common carriage to move product from manufacturing facilities to the retailers and wholesale distributors, this bias is much weaker. Many large manufacturers such as Tyson, International Paper, Kraft, and Archer Daniel Midlands all have sizable fleets that are used in conjunction with For Hire carriers in the execution of their distribution requirements. It is in this portion of the supply chain that this research is most applicable.

To ensure clarity in terms of terminology, I will now forego the use of the standard terms "inbound" and "outbound" freight, since the same load, depending on whether you are shipping or receiving, can be described by either term. Instead, I will utilize the terms "Primary Distribution Network" and "Secondary Distribution Network" as labels that define the echelon in the distribution network that is being discussed.

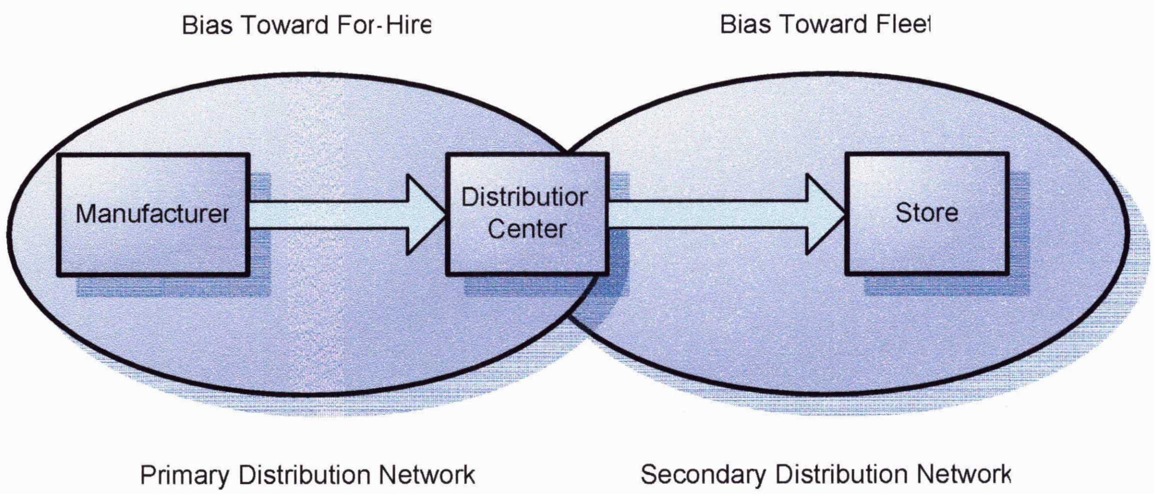

Figure 1 - The Primary and Secondary Distribution Networks

The primary distribution network denotes movements that are from the manufacturer to a distribution center or warehouse. They are often full TL quantities and often cover significant distances.

The secondary distribution network is defined as the movement of product from a distribution center to the customer, typically a retail outlet or store. While big box retailers and grocery chains may have store level deliveries in full or near full TL quantities the majority of shipments in the secondary network are smaller and are supported by pool distributors, LTL and/or parcel carriers as well as private fleets. These fleets cover relatively short distances with high numbers of stops per truck.

It should be noted that there are certainly exceptions to the distribution network described above. For example, direct store delivery is used to bypass the wholesale distribution echelon and service stores directly from the manufacturer. However, the three node network described above is the most commonly employed means to get product from manufacturing to the end market, and will therefore be the model assumed in this research.

Defining Variability

This thesis pays a great deal of attention to the disruptive role that variability plays within a distribution network. Since the term "variability" can have many connotations, I will provide a description of how I will be using the term throughout this research.

For the purposes of this work, "variability" will be defined as changes in volume of shipments moving on a given transportation lane over a period of time, usually a week. The underlying assumption is that the distribution network is relatively fixed in terms of which geographic locations are serviced. In other words, variability is related to the number of loads on a given lane over a specified period of time as opposed to the emergence of new lanes or the removal of existing lanes from the network. The importance of understanding volume variances when creating the strategic transportation policy ensures that the newly formed policy is not based on statistical averages which are often, as we will see, a misleading indicator. Taking into account uncertainty ensures that the strategic transportation policy is developed with an understanding of the typical volume fluctuations that are prevalent in transportation.

Motivation

The need for an efficiently designed and well-implemented transportation policy is becoming critically important for shippers. Trucking de-regulation in the 1980s enabled shippers to negotiate specific contracts with carriers. Well conceived and executed transportation policy gave shippers a new found ability to gain considerable competitive advantage over their peers. The removal of excess truckload capacity in the US transportation market over the past few years has made it even more important.

The convergence of macro-economic, infrastructure, industry specific and supply chain trends that are simultaneously working to increase the net cost a shipper must incur to move product through its network. In the past, network inefficiencies were somewhat shielded due to the market dynamics, which favored shippers over carriers. The threat of losing business in an over-capacitated market was very real to carriers as shippers had the ability to leverage the excess capacity in the market to keep rates low and service levels high.

Since 2003, market dynamics have shifted to favor carriers. Shippers who have built their transportation policy without an eye toward these trends are now being severely penalized, in terms of negotiated rates, service and the amount of capacity allocated to them by carriers. According to the Bureau of Labor Statistics, the price of TL freight has increased 6.7% from November, 2004 to November 2005. Rates are expected to increase an additional 4-6% in 2006 (Standard & Poor's, 2006). Figure 2 shows the convergence of factors that help explain the increase in freight costs shippers are seeing.

Those shippers that are able to best adjust their business to the changing face of the U.S. domestic TL market will be the ones that benefit most significantly while creating a competitive advantage for their business.

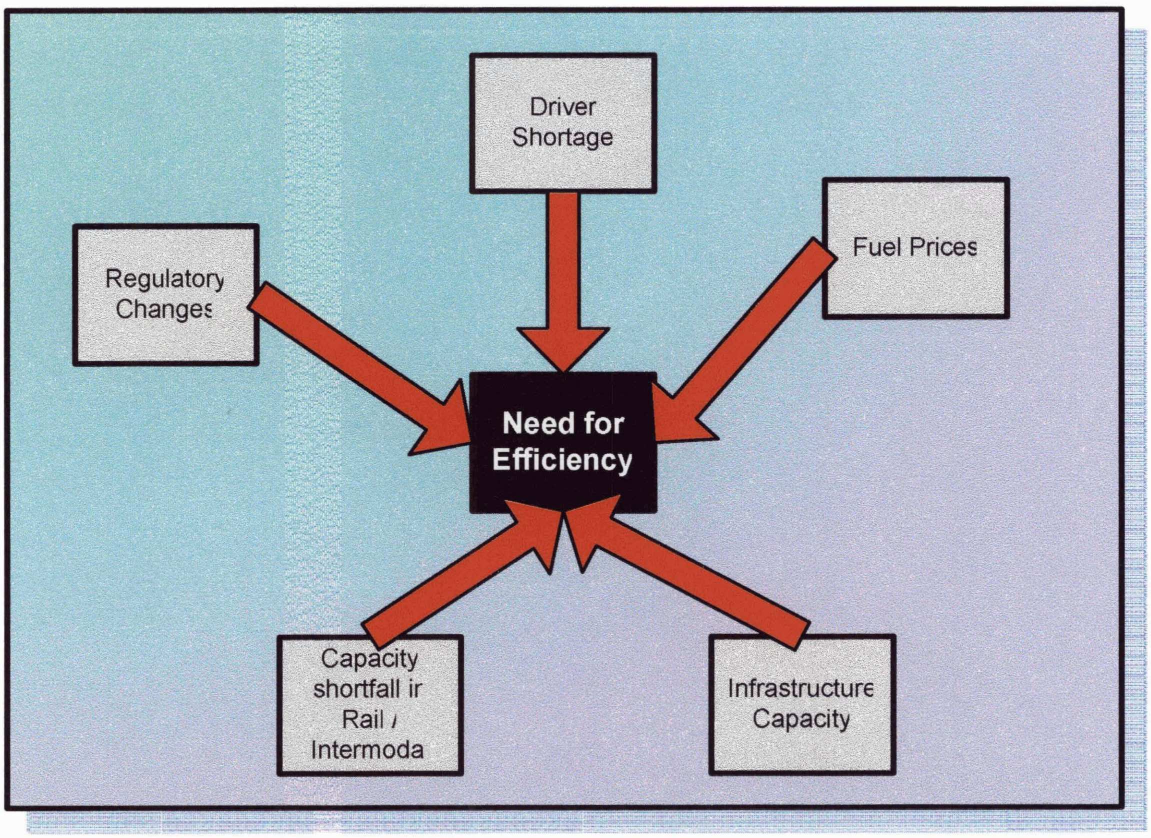

Figure 2 - U.S. Domestic TL Trends

Trends Affecting Domestic TL Transportation

The most obvious trend affecting domestic TL transportation is that of diesel fuel prices. Diesel has risen from $1.50 per gallon in January, 2003 to $2.87 in April of 2006 (Energy Information Administration, 2006). The price of fuel is the second largest expense that national TL carriers incur, after labor, and currently equates to an estimated 16.5% of revenues in 2005, as compared to only 12% or revenues in 2005 (Standard & Poor's, 2006). Fluctuations in fuel prices above and beyond the contracted rates are normally passed along to shippers in terms of fuel surcharges, so while the TL carriers are somewhat shielded from the high price of diesel, the cost does get passed on to the shipper.

Another trend is the ongoing shortage of drivers, especially in the long-haul TL segment. This is further driving up costs as TL carriers look for ways to attract and retain qualified drivers. Carriers are being forced to address driver "quality of life" and compensation issues, both of which will tend to negatively affect the key metrics that carriers and financial analysts use to measure their performance. Shippers are forced to scramble for equipment and are now looking at ways to lock in capacity via long-term contracts and/or the increase usage of private and/or dedicated fleets.

A third trend is the continued introduction of new government regulations that are increasing the burden on both internal and For-Hire transportation fleets. A new Hours of Service (HOS) regulation went into effect in October, 2005 that reduces the number of consecutive hours a driver can work. Additionally, new environmental regulations aimed at reducing sulfur emissions will increase the capital expense of purchasing model year 2007 engines and beyond. These engines require more engine maintenance, are less fuel efficient and require more expensive "clean" diesel.

Finally, the U.S. highway infrastructure has not been able to keep pace with demand. This has caused a 5-fold increase in the total number of hours of delay between 1982 and 2003 for motor vehicle travelers. Total travel delay has increased from 700 million hours to 3.7 billion hours per year (Texas Transportation Institute, 2005). TTI estimates the cost of this congestion to have been $63.1 billion dollars in 2003 for all U.S. motor vehicle travelers.

Without mitigation, these trends will increase a shipper's distribution costs. How much transportation costs increase will vary based on how well a shipper is able to leverage both their internal distribution network as well as their carriers' networks. Fleets, be they internal or For-Hire, will be punished more severely in today's environment than in the past for network inefficiencies. Shippers who are able to create transportation policy in a manner that maximizes utilization of their internal assets while also providing carriers with lane volumes that enhance their network will be able to provide their organizations with reduced operating expenses relative to competitors who are less efficient.

Current Practices

All enterprises translate their anticipated shipping requirements into a transportation policy which is then used as the basis for daily execution. The transportation policy comes in many forms. It can be as simple as a routing guide, which is often a printed out reference guide that tells planners which carrier, mode and service level to use based on criteria such as load size, required service level and origin/destination. Conversely, it can be a complex set of decision variables that include rates, daily or weekly commitments and infrastructure constraints that are used by advanced heuristics within Transportation Management Systems (TMS) to make carrier, mode and through-point selection decisions.

Regardless of how the transportation policy is executed, a fundamental premise of this research is that the tools and methodologies commonly used to create the transportation policy in the first place do not adequately address the needs of organizations whose distribution networks rely on both fleet and common carriage.

In diagnosing the reason for a lack of a solution in this area, the following deficiencies were identified.

1. The transportation policy is created using a deterministic view that does not take into account demand variability.
2. Systems do not take into account the fixed cost of internally owned or leased assets/labor that is charged regardless of use, and therefore do not correctly calculate transportation costs. This biases solutions toward utilizing contract carriers since TMS systems do not understand the sunk cost of a private/dedicated fleet.
3. Systems do not understand the benefits of re-positioning equipment. In many instances, comparing the internal fleet against a common carrier may not be cost effective for an individual load on a given lane, but taken into the totality of the entire transportation plan, the re-positioning of the movement may enable subsequent loads to be moved more cost effectively.
4. Software solution providers have historically concentrated on either the primary distribution network or the secondary distribution network and have tailored their solutions accordingly. Primary distribution network providers such as i2, Oracle,

Manhattan and Manugistics provide solutions that assume For Hire carriers will be used. Descartes, Paragon and UPS Logistics Technologies focus on the secondary network and assume a private/dedicated fleet.

5. The ability of a private/dedicated fleet to generate revenue by moving third party freight is not systematically considered.

Strategic Sourcing of Contract Rates

There are a number of software vendors that are providing tools for contract optimization designed to address the strategic sourcing of For-Hire carriers. Vendors such as Oracle (G-Log), Manhattan Associates, i2 Technologies, and Manugistics/JDA all have solutions that are designed to address transportation contract procurement. These solutions are able to disseminate a shipper's Request for Quote (RFQ) to carriers, collect the carriers' responses and determine the optimal mix of carriers that provides the necessary capacity needed to support the shipper's distribution requirements. In addition to cost optimization, shipper policy constraints, such as core carrier programs are often applied as constraints to the optimization. Examples of common constraints shippers apply include:

- Maximum or minimum number of carriers in the solution
- Maximum or minimum number of carriers per lane
- Carrier fleet size
- Service commitments

While these solutions are able to find the mathematically optimal mix of For-Hire carriers, they do not address distribution networks that utilize both For-Hire and fleet based resources. Additionally, these systems use deterministic approaches. Lane load volumes are converted into weekly averages which are then passed to the carrier who is given little or no information with regard to the variability in demand that will occur over the contracted period. Incumbent carriers who, through previous experience, are aware of the shipper's demand fluctuations often bid higher than carriers that are ignorant of this variability. Carriers that have not worked with the shipper often bid low and win the freight contract, but then have both high turndown ratios in surge periods, as well as underutilized equipment during lulls.

This has a negative impact on the anticipated savings of contract optimization. Shippers are required to stray from their routing guides and utilize more expensive carriers (Harding, 2005) in times of high demand to procure the necessary capacity.

Finally, using a deterministic approach based off of the mean causes carriers' utilization rates to be lower than expected, as half the time the carrier has excess capacity sitting idle while the rest of the time, they do not have enough capacity to meet the shipper's volumes.

Fleet Based Routing Packages

Most shipper-centric fleet based Routing and Scheduling packages focus on the secondary distribution network, moving product from distribution centers to the final customer. Solutions from vendors such as Descartes' and UPS Logistics Technologies are designed to plan multi-stop truckloads in a local or regional delivery area. Due to the perceived need for extremely high levels of customer service, For-Hire carriers are not considered by these packages. All distribution is handled by an internally managed private/dedicated fleet.

Manhattan Associates' Micromap, is a tool that is focused on managing the execution of a long-haul, TL fleet, however it is not useful in the creation of long-term strategic planning in a mixed distribution network. This tool was the output of the work done by Sheffi and Powell, (1988).

Literature Review

The vital role that freight transportation plays in the global economy has encouraged a great deal of research in the use of mathematical as well as qualitative approaches to understand, and ultimately increase the efficiency of many areas of transportation. However, research that addresses the use of private fleets in conjunction with contract carriers is less plentiful. Additionally, utilizing a probalistic approach to data analysis that seeks to incorporate demand variability into the fleet vs. carrier decision process is also an area ripe for research.

Harding (2005) discusses variability in terms of deviations from a shipping enterprise's routing guide due to contract carriers rejecting tendered loads based on variability in supply of transportation capacity or shipper demand. He states that the routing guide, which represents that shipper's transportation policy, is created using deterministic systems and methodologies that do not take into account the inevitable demand or supply variability. Approaches to model the costs of unplanned (loads executed outside the framework of the routing guide) and planned transportation are provided as well as means to mitigate transportation costs. While this approach does provide a great deal of insight into the need to take into account variability when defining policy, it does not provide for the use of an internally managed fleet in the decision making process.

Caplice (1996) discusses how shippers and carriers can share efficiencies by better understanding the role of lane interdependence or economies of scope within a carrier network. By providing loads to carriers that are well aligned with their existing network, the carrier can increase efficiencies, thereby reducing costs for both the carrier and shipper. Caplice provides the means to do this through combinatorial auctions as well as the generation and presentation of continuous tours or "aggregations" in the bidding process. These tours, if serviced by a single carrier, reduce the likelihood that that carrier will incur significant deadhead mileage. By mitigating the risk of empty mileage, the shipper should expect lower transportation costs.

Sheffi and Powell (1988) provides an approach designed for maximizing fleet efficiencies during daily execution. Designed for North American Van Lines (NAVL), the approach describes a framework to maximize profit and customer service in the dispatching process. The goal of the assignment would be to maximize revenue while meeting customer service objectives for the deterministic loads (i.e. Loads that they were currently dispatching), but also use stochastic estimates to position the equipment in an area has a high likelihood of a pickup. By taking into account during the original assignment of the known, deterministic load, NAVL could reduce the risk of having to reposition equipment to a net deficit region. This paper does an excellent job of using probabilities to determine the transactional assignment of a load to equipment, but is more applicable to daily execution than planning and does not provide for the use of contract carriers in the selection process.

Hall and Racer (1995) address the common carrier versus private fleet problem, but do so in the context of private fleet stop-offs versus LTL delivery. They deem stop density and expected weight to be the economic drivers when determining a "break-even" point, at which it becomes too expensive to use the fleet and common carriage should be the preferred method. Ronen (1997) addresses a similar problem, but takes into account varying fleet equipment type. He showed that by minimizing total distance, total cost were 35% higher then by minimizing based on cost.

Johnson and Schneider (1988) published a qualitative study comparing the benefits of private fleet versus contract carriage. Many people predicted that with the deregulation in the trucking industry in 1980, private fleets would become obsolete. The paper discussed why that has not occurred. Some of the reasons cited were changes in regulations that enable private fleets to charge allowances when picking product from suppliers, their use as negotiating tools with contract carriers, their ability to generate revenue and become a profit center through backhauls and third party movements and the continued belief that they provide better customer service and reliability than contract/common carriage.

With the exception of Hall and Race, the research is not developed in a mixed fleet environment, whereby for hire carriers and internally managed fleets compete for the same freight. While Hall and Race do address the common carrier versus fleet decision, the problem they attempt to solve is a trade-off between Less than Truck Load rates and fleet stop-off costs.

I intend to add to the existing literature by developing a framework that develops transportation policy in a Full TL environment by understanding and leveraging the differences in costing between for hire and internally managed fleet assets. Additionally, the framework seeks to provide a level of resilience over standard practices by taking into account lane volume variability in the construction of the transportation policy.

1 – U.S. Truckload Transportation Overview

According to Standard and Poor's survey (S&P, 2006) on commercial transportation released in February, 2006 the market for commercial freight transportation in the United States was $720B in 2004. The trucking market, which is comprised of TL, LTL and private fleets in this estimate, accounted for $671B of the market, which equates to 87.1% of the total U.S. domestic freight market.

The financial contribution of For-Hire TL carriers and private fleets was $606B and was relatively evenly distributed between fleets and For Hire carriers. Private fleets accounted for an estimated $294B of transportation services in the US in 2004 or 44% of the total trucking freight market, while For-hire TL carriers represented $312B of the market.

In order to understand how to create a well conceived transportation policy, it is imperative that the reader have a base understanding of how private, dedicated and For-Hire carriers differ in terms of cost structures and services provided. This section will provide a high level overview of each of these different types of transportation.

Trends and how they affect domestic TL transportation will also be discussed so as to provide the reader a clear understanding as to the short and long term challenges that shippers and carriers face. The intent of describing these trends is to ensure that the transportation policy is developed in a manner that addresses the challenges and does so in a manner that maximizes the competitive advantage of the shipper over its peers.

1.1 - Private Fleets

A private fleet is owned and operated by the shipping entity. The principle objective of the fleet is to support the shipper's internal distribution requirements. The shipper leases or owns the physical assets such as tractors, trailers and/or straight trucks. The drivers are normally employees of the company. Private fleets can attain common carrier authority and act as a For-Hire carrier, moving freight for other shippers. This is a point that will be discussed and leveraged later in this thesis.

The 10 largest private fleets (Transport Topics, 2005) are owned and operated by the following companies.

Company
1. Sysco Corporation
2. Wal-Mart Stores
3. Ahold USA/US Foodservice
4. Tyson Foods
5. McLane Co.
6. Schlumbergr Ltd.
7. BJ Services Co.
8. Kroger Co.
9. Pilgrim's Pride Corporation
10. Albertsons Inc.

Table 1 - Top 10 Private Fleets in the U.S. as of 2005

Private fleets are prevalent in industries such as grocery (Albertsons, Kroger) and foodservice (Sysco Foods, Ahold/US Foodservice). Of the top 100 private fleets, over 50% are involved in the distribution of food products. The wholesale foodservice industry is one of the largest users of private fleets in the country. This industry supports the delivery of food products from the wholesale food distributors to their customers. These companies purchase food products from manufacturers and provide local distribution to restaurants, bars, hotels and other businesses that are too small to warrant receiving orders directly from the manufacturer.

Private fleets in this environment typically depart and return to the same distribution center each day, and deliver orders to customers who place orders the previous day. Order quantities vary from under 100lbs up to a few thousand lbs of product per delivery. The wholesale foodservice industry averages between 15 and 25 customer deliveries per truck per day, often with customers expecting consistent delivery times and pre-established delivery days.

Customer service is the preeminent driver of these types of fleets trumping in importance optimal routing sequences and vehicle utilization. These short-haul runs often are executed using 28' "pup" trailers and straight trucks (trucks that have the storage area mounted to the chasis), as opposed to the larger (48' and 53') tractor trailers typical of long-haul transportation. Because of the nature of this business, routes are typically constrained more by delivery time than physical space on the vehicle. Optimization is typically used during tactical planning in order to create driver territories and establish route sequences, but due to the expectation of high service levels, static route planning is often used for daily execution.

For-Hire carriers are not considered in these operations due to the need for extremely high level of customer service. Fleets are also much more cost effective than for-hire carriers in these operations due to the high density of low volume stops. This allows a single fleet truck to support the delivery requirements of many of the wholesale distributor's customers. No cost-effective For-hire alternative exists for this type of network, so peddle route fleets are almost exclusively used in this environment.

Of more applicability to this research are shippers that use private fleets for longer haul, full truckload transportation. This is common in many industries including grocery retail (Albertsons, HEB), big-box retail (WalMart) and food/consumer product goods (CPG) manufacturers. In this segment, for-hire carriers are often times used to supplement fleet capacity. Typically, for-hire carriers are more efficient due to advantages they hold in terms of economies of scale and scope, however, many shippers still view fleets as an important component in their overall transportation strategy. Common reasons cited for having a fleet include:

- Better perceived service to their customers. Fleet drivers are viewed as important assets in maintaining a strong shipper/customer relationship.
- Fleet drivers can be requested to perform special services during the delivery that For-Hire carriers will not do or would do only for an additional charge.
- More leverage with contract carriers during rate negotiations by sending a message to the carrier that it can be replaced by an internally managed fleet
- Marketing advantages of having the shipper's name on the trailer, thereby acting as a rolling billboard for the company.
- Provides assurance of freight capacity times of tight capacity, such as exist in the current environment.
- More control over transportation operations

Michelle Livingstone, Sr. Vice President of Transportation at C&S Wholesale distributors put it this way when asked why she felt fleets were important part of a comprehensive transportation strategy.

"In order to ensure there is sufficient capacity at all times and to quickly respond to customer requirements, I believe the best transportation network is comprised of a private fleet, dedicated fleets and one-way carriers. Our fleet is larger than many common carriers so we enjoy many of the same economies of scale that a common carrier would." (Livingstone, 2006)

1.2 - Dedicated fleets

Unlike private fleets, dedicated fleets are not owned by the shipper, but are provided on an exclusive basis to the shipper for a contractually specified period of time. Most large TL carriers like Schneider National, JB Hunt, Swift and Werner have active and growing dedicated fleet businesses. The advantages of a dedicated fleet over a private fleet is that a dedicated fleet does not require a large capital expenditures outlay as is required when expanding the capacity of a shipper's private fleet. Dedicated fleets provide the advantages of guaranteed capacity in constrained markets and increased control over the asset and its usage. This has become an important advantage as shippers compete with each other for the available TL capacity on the market (Bradley, 2005)

A director of one of the world's largest global third party logistics companies puts it this way when describing the benefits dedicated fleets over contract carriers.

"Typically most dedicated fleets would be used over contract carriers due to service requirements...possibly specialized equipment and/or specific services that the drivers perform as part of their overall duties are simply difficult (if not impossible in some cases) to replicate with contract carriers...there is benefit of having the drivers and equipment under your direct control...equipment can also be used for marketing

and/or advertising purposes that simply cannot be done with contract carriers." (Mulqueen, 2006)

Like private fleets, shippers will incur dedicated fleet costs regardless of whether the assets are used since a large component of the cost is fixed versus load based, as in the case of contract carriage. Idle fleet assets and excessive dwell time are especially costly for dedicated as well as private fleets. Additionally, the variable component of dedicated contracts typically includes per mile charges that are incurred regardless of whether the vehicle is loaded, thereby penalizing shippers that have high levels of inefficiencies, as defined by deadhead miles in their networks.

Dedicated and private fleets are both most effective when the shipper has the ability to maximize equipment utilization by minimizing deadhead distance and dwell time as well as by utilizing the assets on shorter-distance runs that incur high minimum fees when executed by for-hire carriers.

While there are key differences between private and dedicated fleets, this thesis will view both as fundamentally the same. Both modes have significant sunk costs and are perishable in nature in the sense that if the capacity is not used, it is lost. This contrasts with For-Hire carriers, which are paid only when used to execute the movement of a load.

1.2.1 - Private Fleet and Third Party Freight

A benefit that some companies take advantage of with regard to their private/dedicated fleets is the ability to generate revenue by moving freight for other shipping entities. As the TL capacity in the United States continues to tighten and rates are driven up, it is becoming economically compelling for private fleets that have significant empty miles built into their network to acquire common carrier authority and move other shipper's freight.

This enables the fleet to generate revenue and turn into a profit center for the shipping organization. Fifty-six percent of private fleets operate with common carrier authority today (Terreri, 2006), although it is not known what percentage of these fleets are actually moving third party freight since there are benefits aside from generating revenue that drive a fleet to attain common carrier authority.

Fleets gain visibility into third party freight through existing relationships with trading partners, transportation brokers such as CH Robinson and/or the use of on-line transportation exchanges that seek to bring buyers and sellers of transportation services together. In theory, by enabling sellers of transportation services to place available capacity on the open market, synergies between shippers and carriers will be discovered and the overall network will become more efficient. This led to a large number of public and private transportation exchanges emerging in the 1990s, most of whom are no longer in business. In reality, exchanges and the spot market are often used by shippers only as a last resort when all other options have been exhausted. Shippers have proved leery to give freight on a regular basis to carriers to whom they do not have an existing relationship. However, as TL capacity continues to tighten, utilizing exchanges and brokers to procure and sell transport capacity may become more necessary to maximize efficiencies in the U.S. domestic TL network.

1.3 - For-Hire TL Carriers

There were approximately 45,000 for-hire truckload carriers in the United States in 2004, of which approximately 30,000 (67%) had annual revenues of under $1M (Standard & Poor's, 2006). The large number of TL carriers is due to the low barriers of entry required to begin hauling freight TL freight, which unlike the parcel and LTL segments, does not require capital investment in physical infrastructure such as consolidation terminals. The largest TL carrier, Schneider National, had TL revenues of approximately $3.1B or only 1% of the market in 2005.

A For-Hire TL carrier is contracted by outside organizations to move freight at a pre-determined rate and operate in environments where loads are greater than 10,000 lbs, which is the approximate breakpoint where the variable nature of LTL costs begin to exceed the fixed nature of TL costs. These carriers pick up freight at the origination point and move it to the final destination without any intermediate loading and unloading of the shipment, although the shipper can contract for the TL carrier to perform multiple pickups and/or deliveries under the same bill of lading. This is markedly different from the LTL and parcel transportation network models, which utilize hub and spoke systems that require multiple transfer points to move product from the origin to the ultimate destination.

The primary benefit of a hub and spoke network is that it enables consolidation of shipments going between terminals. This benefit is not recognizable in a TL environment, since the vehicle is, theoretically, already fully utilized and injecting a full TL into a hub and spoke network would simply add additional transit time and handling expense to the process.

Aside from rates, TL contracts with shippers define the "rules of engagement" in terms of how each party will operate within the relationship. For instance, the contract may specify that the carrier must be able to send and receive the standard TL Electronic Data Interchange (EDI) transactions for load tenders, load accept/reject, shipment status and freight invoice.

Additionally capacity commitments are often specified within TL contracts. From the shipper's perspective, capacity commitments require the carrier to cover a certain number of loads on a given lane over a specified period of time. This is often done by requesting that the carrier agree to haul a set percentage of total load volume on a each lane. This, in theory, provides the shipper the capacity needed to manage the weekly variations in load volumes that a fixed volume commitment would not support.

Origin	*Destination*	*Commitment*
Boston, MA	Atlanta, GA	75%

Table 2 - Contracted Lane Allocations

Table 2 shows a simple example of a capacity commitment that specifies that regardless of the shipper's load volume, the carrier is responsible for moving 75% of the freight on this lane. Whether there are 10 or 20 loads, in this scheme, the carrier is responsible for flexing capacity to meet the contractual commitment. Often times these commitments are made without any sense of the volume variability on the lane.

Similarly, the carrier may request a certain amount of load volume over a set period of time that if agreed to, will earn the shipper a lower rate. In this example, the carrier may give a shipper a better rate if the shipper commits to 10 loads per week as opposed to only 5. This example of economies of scale in carrier negotiation is a less powerful motivator of rates than network efficiencies seen with economies of scope, as we shall see discuss later in this research.

1.3.1 - Carrier's Ability to Decline Loads

One important facet of TL transportation that needs to be recognized is that unlike LTL or parcel carriers, For-Hire TL carriers will often reject undesirable loads; even those loads under contract. This occurs if the carrier does not have available capacity or, as carriers get more technologically savvy the load is deemed operationally unprofitable given the current location and status of the carrier's assets. The frequency of carriers turning down loads has become more prevalent in recent years as US domestic TL capacity has tightened and the carriers have begun to exert their new found power in the buyer/seller relationship. Harding (2005) showed that the cost of a turndown was estimated to be between 2% and 7% of the freight spend. In this analysis, over 25% of tendered loads under study were rejected by the primary carrier. The effect of declined freight is discussed extensively by Harding (2005).

1.3.2 - TL Rates

TL rates are negotiated as one-way, door to door rates. The shipper is charged for a given load based on the origin and destination pair. A key difference between For-Hire carriers and dedicated/private fleet is that the shipper is not responsible for the mileage or time required to get the TL carrier's equipment and driver to the load origin point, nor is the shipper responsible for the equipment moving from the contracted destination point to the carrier's next load origination point, unless under the special condition of a continuous move, as discussed later. This is an important point, since carriers will negotiate higher rates for shipper lanes that do not fit well into their existing network, but conversely, lanes that add to a carrier's overall vehicle utilization, as defined by the ratio of $\frac{LoadedMiles}{TotalMiles}$ will be priced more aggressively.

An efficient network that minimizes deadhead distance and dwell time is critical to a For-Hire carrier just as it is to the shipper when it manages its internal private or dedicated fleet. A fundamental metric that investors look at to determine the efficiency of a TL operation is the operating ratio. This is defined as:

$$OperatingRatio = \frac{OperatingExpenses}{Gross\,\mathrm{Re}venue}$$

A low operating ratio (.85 or lower) typically means that the fleet is being well utilized. Loads tendered to a carrier that require excessive, unpaid re-positioning of assets to the load origin point or from the load destination point will normally elicit higher rate quotes or the carrier simply declining the load in order to keep the operating ratio in-line with corporate objectives.

1.3.3 - TL Rate Structure

While contracts between shippers and For-Hire TL carriers utilize a myriad of pricing schemes, a typical contract TL tariff is described below.

Tariff Attribute	*Description*	*Example*
Origin Zone	Origin zones that specify the origination point of a	Boston, MA

	load. They are standardized geographic areas such as postal codes, cities or states/provinces.	
Destination Zone	Destination zones specify the destination of a load. They follow the same guidelines as the origin zone.	Atlanta, GA
Equipment Class	Equipment class is used to specify the type of equipment required for a given load. Examples of equipment classes are Dry Vans, Refrigerated Vans and Flat beds.	Dry Van
Team/Single	This attribute relates to the number of drivers that are available to operate the truck. Teams have the advantage of being able to operate 24x7, while single drivers must adhere to Hours of Service rules, which mandate break periods.	Single
Cost / Mile	This is the cost per mile. Most carriers and shippers use mileage calculation tools to such as ALK's PC Miler solution to ensure that mileage is appropriately and consistently calculated across enterprises.	$1.82
Stop-Off Charges	TL Carriers quote rates for a single pickup and delivery. However, shippers have the ability to request that carriers pickup and/or deliver to multiple locations under the same bill of lading.	$120/stopoff
Other Accessorials	Other accessorials such as detention charges, lumper fees, demurrage and out of route miles may also be applied. Many of these charges are not known at the time of planning, and are only known to the shipper upon receipt of the freight bill.	

Table 3 - Attributes of a For-Hire Carrier Rate

1.3.4 - Continuous Move Rates

Continuous move rates are TL rates that provide discounts off of the standard, one-way rates. They are negotiated when a shipper is able to provide to the carrier subsequent loads that originate within a specified radial distance and agreed-to time period of a preceding load, thereby increasing the productivity of the carrier's fleet.

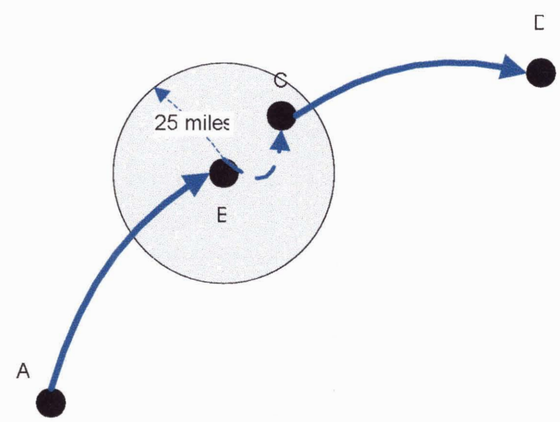

Figure 3 - Continuous Move

Figure 4 represents a load originating at point A destined to point B. The shipper then provides the carrier a subsequent load, from point C to point D. If the load is within the agreed upon temporal and spatial limitations, the carrier will provide the shipper a discount for one or both of the loads. The discount is normally negotiated to be a percentage off of the standard one-way rate. Additionally, the dead-head distance between loads is normally heavily discounted or waived.

While continuous move rates are less commonly employed relative to point to point rates TL carriers typically quote, their very existence shows the potential value of a shipper understanding a carrier's financial drivers in terms of lane interdependency when contracting and executing freight.

1.3.5 – The Cherry Picking Problem

An issue was raised when discussing this solution with large shippers about the potential impact on For-Hire carrier rates if the best lanes, i.e. those with high volume and low variability, are reserved for the fleet. The issue, as they see if, is that by "cherry-picking" the t lanes for the fleet, the savings a shipper will see in terms of increased fleet utilization will be off-set by the increase in cost associated with providing the carrier the least desirable loads.

While this concern must not be discounted, it should be noted that lane variability and volume are not the only criteria used in tour formation. The principle objective in the tour assignment process is to maximize cost savings when comparing For-Hire and fleet resources. Assuming that the underlying premise of this thesis regarding economies of scope is valid, For-Hire carrier rate discrepancies are in large part driven by how well the lane fits into the carrier's network. Lanes that fit well within a carrier's network will be discounted by the carrier while those that do not fit well will be charged a premium. With the objective function attempting to maximize savings, the bias, with all else equal, will be to use the internal fleet on lanes that, through pricing, we can infer are less desirable to the carrier.

1.4 – Shipper and Carrier Challenges

A combination of strong economic growth, changes to regulatory policies, labor shortages and rising fuel costs have created a severe shortage of transportation capacity in the United States trucking market. The capacity shortfall has changed the dynamics of the relationship between shipper and carrier. Carriers today enjoy a "sellers market," providing them with much more leverage than they have historically had when dealing with shippers. Carriers can be more selective in determining which shipper's freight to haul as well as much more aggressive in terms of pushing forward rate increases to their customers. They are also more prone to charge shipper's for poorly run operations that result in increased carrier wait time for pickups or deliveries, incorrect documentation or poorly staged freight (Bradley, 2005).

Outlined below are the economic, regulatory and demographic realities that shippers and the long-haul trucking industry face. Understanding these trends as well as their underlying ramifications is essential to the creation of long-term transportation policy that best mitigates these potentially disruptive trends and occurrences.

1.4.1 – Driver Shortage

An acute driver shortage that started in the late 1990s shows no signs of abating. A study commissioned by the American Trucking Association and conducted by Global Insight, Inc estimates that there is currently a shortage of 20,000 long-haul drivers in the United States. That shortage, assuming no large-scale market changes intercede, will grow to over 110,000 by 2014 as the industry fights the reality of a demographic shift in the market. The trends cited within the report include:

- ***Less New Workers Entering U.S. Labor Force*** - The growth of the US labor force will slow from its current level of 1.4% annually to .5% by 2012.
- ***The primary demographic of long-haul truckload drivers is shrinking*** - The white, male population between the ages of 35-54 will decline by 3 million people by 2014. This group currently makes up over 50% of the total number of long-haul drivers.

- ***Driver Retirements Looming*** – One fifth of long-haul drivers today are over 55 years of age.
- ***Strong Competition for labor between Construction and Driving*** – Workers often flow between construction and truck driving. In the 1990s, average long-haul truck drivers earned 6-7% higher wages than construction workers. The wage premium drivers historically have enjoyed has been eliminated, with construction workers now earning slightly more than drivers.
- ***Competition for drivers between LTL/Parcel and TL*** – LTL and small package drivers are, on average, better compensated than TL drivers. The average LTL driver earned over $42,000 in 2003, vs. just over $36,000 for TL drivers (Standard and Poor's, 2006). Better compensation, along with regional routes that provide more time at home make these segments more attractive to drivers vis-à-vis the long-haul TL segment.

Aside from lower wages, long-haul truck drivers frequently cite "quality of life" issues as fundamental reasons for choosing to leave long-haul trucking. These issues include being away for extended periods of time as well as irregular and unpredictable route schedules. In order to mitigate the loss of drivers due to this issue, TL carriers are now looking to reconfigure their networks to better address these issues. The trade-off under consideration is the cost of more empty miles to ensure driver's return home more frequently vs. the costs of attracting and hiring new drivers due to very high levels of driver attrition.

The driver shortage is also having a noticeable affect on the ability of large TL carriers to grow revenues. Below are statements from two of the largest publicly traded carriers - JB Hunt and Werner Enterprises - regarding the current driver shortage.

"Driver and independent contractor availability continues to be a serious concern for the segment, as well as the industry. We continue to see no signs of fundamental improvement in driver or independent contractor availability for the foreseeable future. Therefore, we do not anticipate significant capacity additions in the truckload market in the near term." (JB Hunt, 2006)

"The driver recruiting and retention market remains more challenging than ever. The supply of qualified truck drivers continues to be constrained due to alternative jobs to truck driving that are available in today's economy. The Company continues to focus on driver quality of life issues such as developing more driving jobs with more frequent home time, providing drivers with newer trucks, and maximizing mileage productivity within the federal hours of service (HOS) regulations"(Werner Enterprises, 2006)

1.4.2 – Regulatory Changes

New regulatory policies aimed at addressing safety and environment concerns have been either recently implemented or approved. These regulations are going to further reduce the overall productivity of carriers and fleets.

Hours of Service Changes

A new Hours of Service (HOS) regulation went into effect on October 1, 2005. Implemented by the Federal Motor Carrier Safety Administration (FMCSA) in the United States, the aim of the revamped regulation is to address safety concerns related to driver fatigue. The net effect of the regulation is that it reduces the total time a driver is able to be on-duty from 15 to 14 hours. Certain activities, such as break times, that were not counted against driver's hours under the old guidelines are now being counted, which effectively reduces on-duty time even further.

The regulation also requires that drivers now must be off-duty for 10 hours instead of 8 before going back on duty. Taken together, these regulations will have a net adverse affect on overall productivity. Wal-Mart has estimated that the regulations will cost the company an additional $24M / year due to the need to acquire additional equipment and hire more drivers to account for the anticipated decrease in transport productivity.

1.4.3 – Updated Emission Standards

New emission standards put forth by the United States Environmental Protection Agency will impact all 2007 model and beyond heavy-duty engines. While more eco-friendly, these engines require a cleaner, but more expensive type of diesel fuel. The EPA has estimated that the fuel will cost an additional 4.5-5 cents per gallon. The engines are also expected to be less fuel-efficient and require more maintenance than engines built prior to 2007. Finally, the engines will cost more than the 2006 and prior year models. The increased cost estimates vary from approximately $2000 to $5000 per engine.

1.4.4 - Fuel Prices

Fuel prices have risen to their highest levels in history in the past few years in the United States. Between January, 2003 and April 2006, average US diesel fuel prices have risen 91%. Fuel prices now comprise the second largest portion of operating expenses of TL carriers, just behind driver compensation. According to their respective annual reports, fuel charges were 22% of the operating expense for Swift Transport and 21% for Werner in 2005.

To mitigate price fluctuations, carriers are instituting fuel surcharges. Carrier fuel surcharges are based on the department of Energy's U.S. National Average Fuel Index and typically update on a weekly basis. They provide a means to ensure that the fluctuations in the price of diesel are not born solely by the carrier, but are passed through to the shipper.

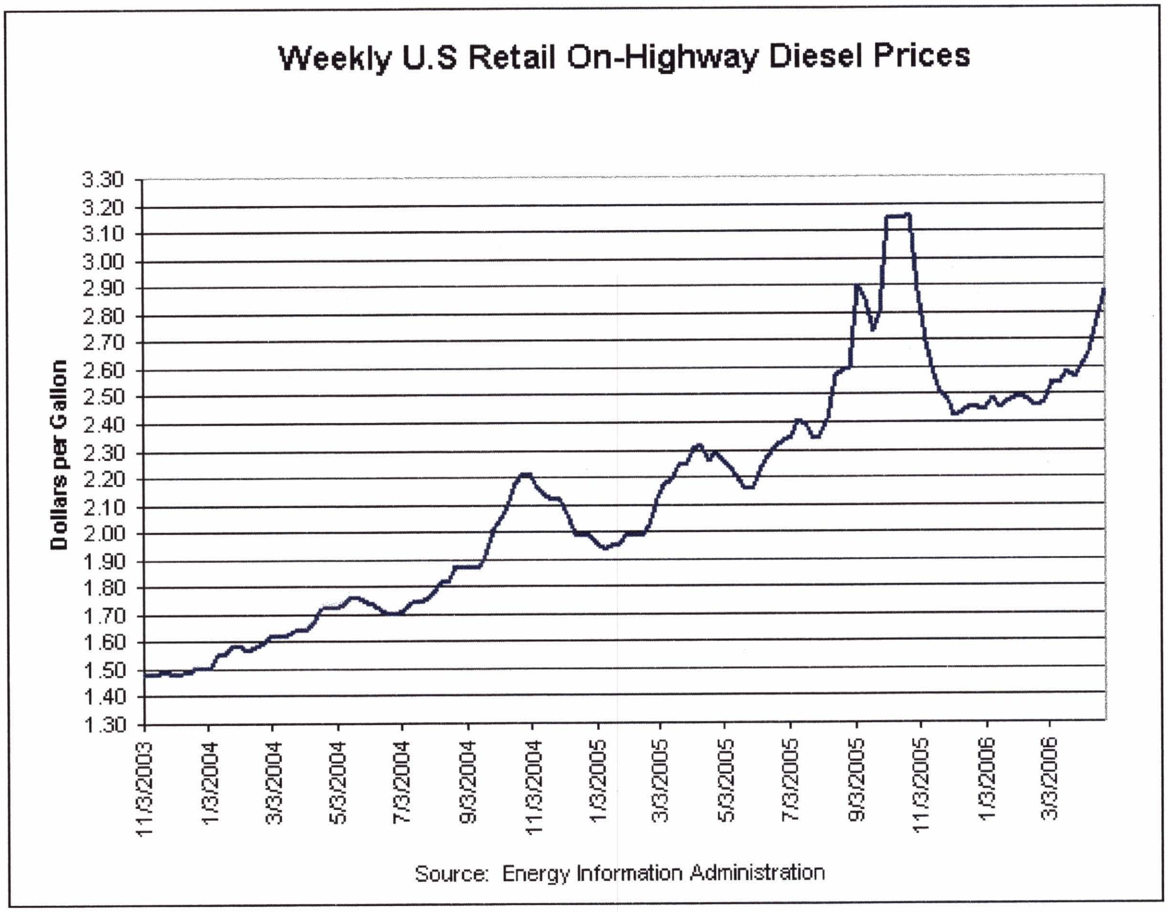

Figure 4 - Diesel Fuel Prices

1.5 – Transportation Exchanges

Transportation exchanges are on-line communities designed to leverage the power of the internet by bringing together buyers and sellers of transportation services. Examples of transportation exchanges are Transcore and Get Loaded. In these exchanges, shippers have the ability to place loads up for auction. Similarly, carriers can place excess capacity up for sale allowing shippers to bid.

In this research, I suggest that exchanges can be used by shippers as the means to make their fleet's excess capacity available on the open market, thereby turning their private fleet into a profit center and reducing empty miles. This is a somewhat contentious issue with many fleet managers who are unwilling to risk utilizing their fleet for anything aside from the fleet's primary purpose, which is to support the shipper's own distribution requirements.

In a focused interview with 37 shippers in August, 2004 14 of 37 (38%) shippers stated their willingness to haul another shipper's freight if the alternative was significant empty miles. (UPS Logistics Technologies, 2004) While not statistically significant, the questionnaire does show that under the right conditions, some shippers are willing and able to haul third party freight.

2 - Data Characterization

In this chapter, I will introduce the reader to the relevant characteristics of a typical distribution network to which this research applies. The intent is to provide insights into two specific areas. First, an analysis of how this shipper was charged for freight will look to identify characteristics that drive the variability in rates provided by carriers. Second, the research will delve into lane variability seen within the sample data set and discuss how this information can be effectively used when formulating transportation policy.

2.1 Economies of Scale and Scope

Economies of scale - as it relates to transportation in general and this research in particular - assumes that the greater the load volume on any given lane, the lower the average cost per load on that lane will be. For instance, in a network where there are strong economies of scale, a shipper with 100 loads per week on lane A to B will be given a much better rate than a shipper with 10 loads per week on the same lane. Using the strictest interpretation, economies of scale assumes complete lane independence in terms of carrier costing, whereby freight rates are determined solely based on volume and not on how well the loads integrate with the existing transportation network.

Economies of scope are seen when rate reductions are provided not based on lane volumes, but how well the freight lanes in question integrate with the rest of the carrier or fleet's network. Economies of scope assume that the driving force behind rate variability is how well lane volumes interact with the existing network.

A fundamental premise of this research is that economies of scope are a much stronger force than economies of scale in domestic TL transportation. This is discussed in the following analysis however a lack of understanding of how this shipper's lane volume impacts the carrier's network makes proving this premise impossible. Regardless, the assumption that this research was written under assumes that this premise is valid and holds true for For-hire carriers as well as internally managed private/dedicated fleets.

2.2 - Data Overview

The data analyzed is from a mid-to-large sized consumer package goods manufacturer distributing food products in the United States. The data consist of approximately 230,000 loads executed over a 51 week period. The loads emanate from half a dozen origin points and represent approximately $200 million dollars of annual TL freight spend. The freight is transported just over 99 million miles with the average load traveling just over 400 miles. All loads are full truckloads (TL) originating and terminating in the United States. There are a small number of load origin points distributing to a large number of destination points, which is indicative of an outbound distribution network, as perceived from the manufacturer's perspective.

The loads cover 1838 distinct origin and destination points or lanes. Lanes are defined based on postal codes, whereby all loads that share the same origin and are destined to the same 3 digit zip code are aggregated together. For example, a load originating out of the Georgia facility and terminating in postal code 21030 would be aggregated with another load originating from the same Georgia facility but terminating in postal code 21099, since the destinations share the first three digits of their postal codes, which in the United States, specifies a specific region of the country.

The load breakdown by facility is detailed in table 4.

Load Origin	Total Loads	Avg Load Distance	Total Cost	Cost / Mile	Total Distance	Avg Cost/Load
DC1	52308	418	$48,843,077	$2.23	21,890,237	$933.76
DC2	50015	357	$44,589,104	$2.50	17,856,300	$891.51
DC3	47487	572	$45,475,009	$1.67	27,171,362	$957.63
DC4	38913	496	$36,860,558	$1.91	19,283,847	$947.26
DC5	41978	286	$21,595,682	$1.80	12,011,566	$514.45
DC6	604	704	$601,953	$1.42	424,988	$996.61
DC7	1440	748	$2,186,763	$2.03	1,076,700	$1,518.59
	232,745	*428*	*$200,152,146*	*$2.01*	*99,715,000*	*$859.96*

Table 4 - Loads by Origin Point

2.2.1 – Network Costs

The total freight spend over the 51 week period was approximately $200 million. Just over 21% of the 1838 lanes accounted for 80% of the freight cost, with the bottom 50% of the lanes in terms of freight cost accounting for just 4.83% of the total transportation spend (Figure 6).

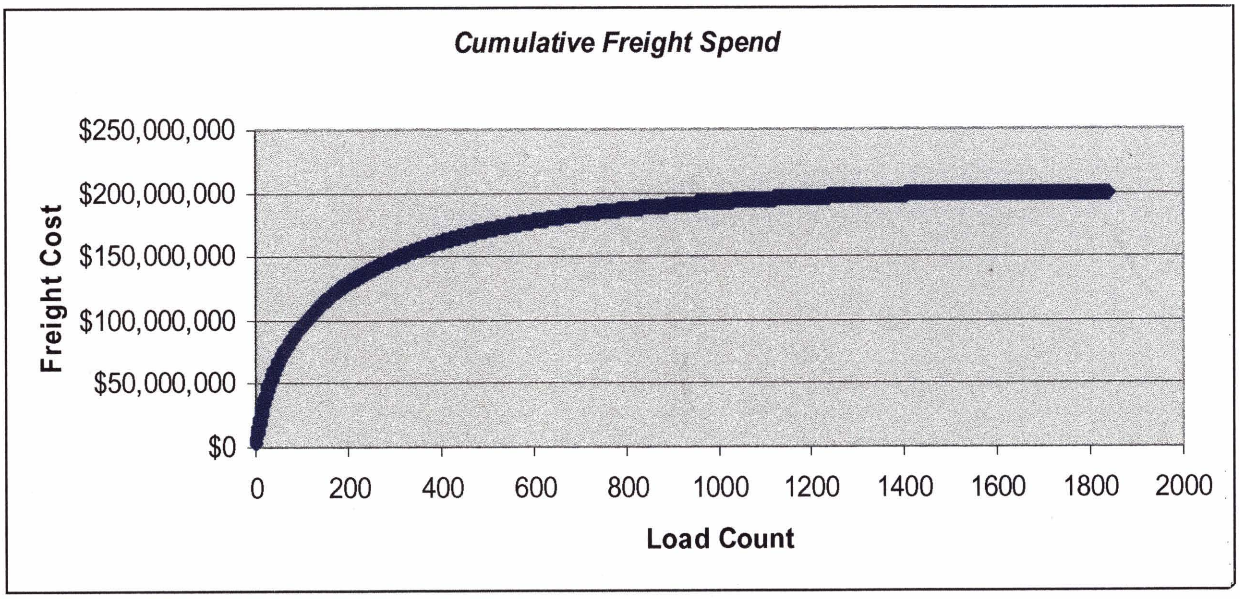

Figure 5 - Annual Freight Spend

2.2.2 – Freight Cost and Distance

While a strong correlation of .775 was found between the average lane cost and the distance of the lane for the data analyzed, transportation costs did vary significantly between origin points. For the 5 manufacturing facilities that make up the bulk of the loads studied, average per mile rates range from $1.67 for loads originating at the Texas facility up to $2.50 for loads originating in North Carolina. When removing the distance variability between the load origination points, there still remained a great deal of variability in terms of average cost / mile.

Figure 7 shows that rates varied from a low of $1.35 / mile for loads originating in Oregon up to $2.80 for loads originating in North Carolina when limiting the analysis to only loads between 500 and 600 miles. While distance is an important driver of costs, it is apparent that there are other factors that are driving the rate disparity across manufacturing facilities.

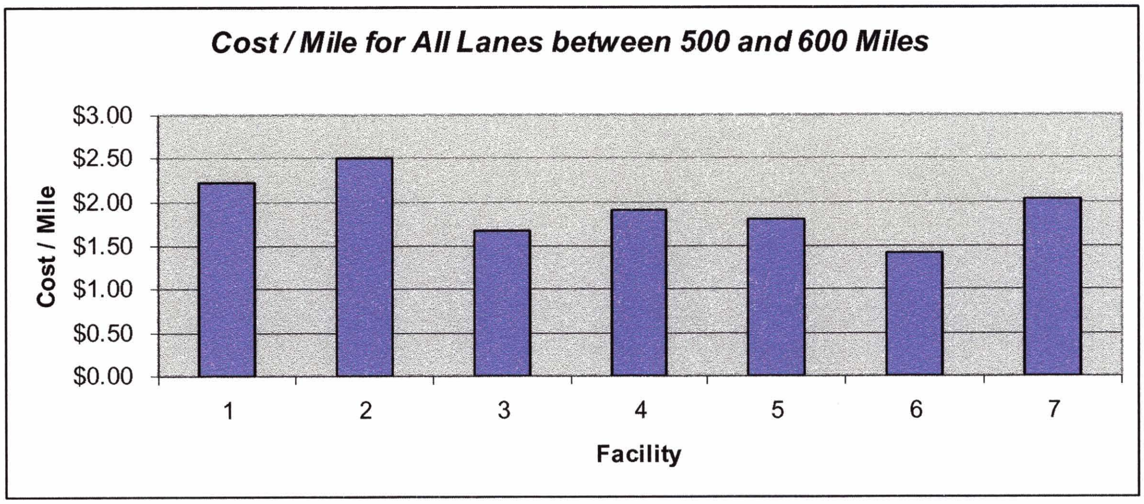

Figure 6 - Cost per Mile by Facility

Even lanes originating at the same facility and traveling relatively similar distances have a good deal of variability in average cost per mile. For instance, 38 lanes comprising 10,072 loads originate at the Georgia facility and require transport of between 400 and 600 miles. The average rate in this group was $2.07/mile however the spread between the lowest rate and highest rate was $1.51/mile to $3.57/mile with a standard deviation of .60¢. Again, this shows that while distance is a factor in rate variability, other factors are also at play. This is shown in figure 8.

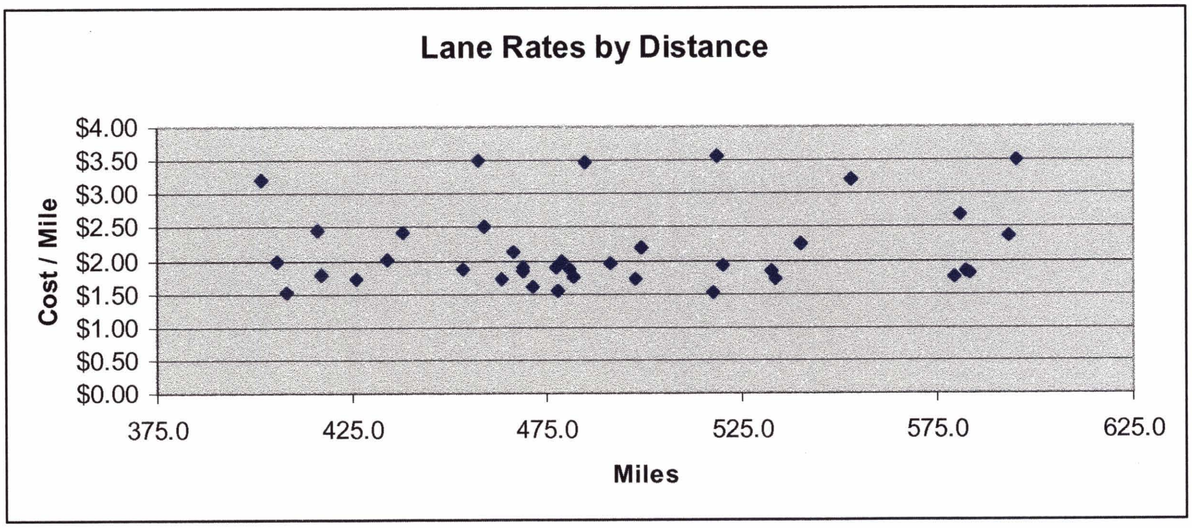

Figure 7 - TL Rate Variations

In order to help understand why these types of rate discrepancies exist, an analysis was done to determine the effects of lane volume on rates. To determine the strength of economies of scale, a correlation between rates and lane volumes was calculated for all lanes between 400 and 600 miles. A strong negative correlation between high lane volumes and low rates would indicate that economies of scale are a powerful force in TL transportation pricing, however, the correlation, as anticipated came back relatively weak (-.118 correlation).

For the entire data set, the correlation between lane volume and cost / mile was .067, which can be interpreted to mean that there was little in the way of relationship between high lane volumes and lower rates. What cannot be rejected is the possibility that the entire volume of the shipper's freight did enable it to receive lower rates across the board when it negotiated contracts with its carriers. While the economies of scale were not apparent when looking within the shipper's network, the total lane volume may have been discounted by carriers in rate negotiations based on the entirety of volume bid open when compared to other, smaller shippers. This cannot be ascertained from the available data.

To fully determine the affects of economies of scope on TL pricing requires an in-depth understanding of each carrier's internal transportation network. Since this is not known, and normally will not be known by a shipper when formulating transportation policy, we can only look to see if lower rates are provided for lanes that, within the shipper network, are synergistic. To investigate this, an analysis was done to see if lanes that moved from a manufacturing facility to a region that contained another manufacturing facility were discounted when compared to similar loads.

Rates for lanes that terminated in the 3 digit postal code which corresponded to one of the manufacturing locations were compared against rates for lanes that originated from the same facility and traveled within 5% of the distance of the first load, but did not terminate in the same region as one of the shipper's manufacturing sites. The data show for most lanes, there are significant discounts for lanes that terminate in a known load origination point, in this case, the shipper's manufacturing site. This is shown in table 5 and while not conclusive, does support the premise that economies of scope are a driving factor in carrier pricing that are responsible for a portion of the variance in rates we see within a network. This is an important factor that shippers need to recognize when creating transportation policy.

Lane	Distance	Rate to Manufacturer Region	Rate to non-Manufacturer Region	Percentage Discount
NC to GA	497	$1.92	$2.62	26.72%
NC to OH	449	$1.76	$2.66	33.83%
GA to OH	653	$1.76	$2.57	31.52%
GA to NC	482	$1.60	$2.04	21.57%
WI to NC	785	$1.46	$1.73	15.61%
WI to OH	361	$1.83	$1.43	-27.97%
TX to GA	855	$1.43	$1.77	19.21%
TX to CA	1383	$1.34	$1.63	17.79%

Table 5 - Economies of Scope

2.3 Lane Volumes and Variability

Every shipper must support some level of volume variability in their transportation plan. Seasonality, product promotions, unpredictable customer ordering patterns, new product roll-outs and infrastructure changes all contribute to variability that the transportation organization must be able to address. By understanding normal demand variability, transportation policy can be crafted in a manner that seeks to ensure relatively high levels of fleet utilization for demand that is more certain and increased reliance on For-Hire carriers will demand that is less certain.

2.3.1 - Load Frequency by Week

Over the 51 weeks analyzed, weekly load volumes ranged from a minimum of 3072 loads to a maximum of 5975. The average weekly load volume was 4,563 with a standard deviation of 625. The coefficient of variation, defined as $\frac{\sigma}{\mu}$, = .137.

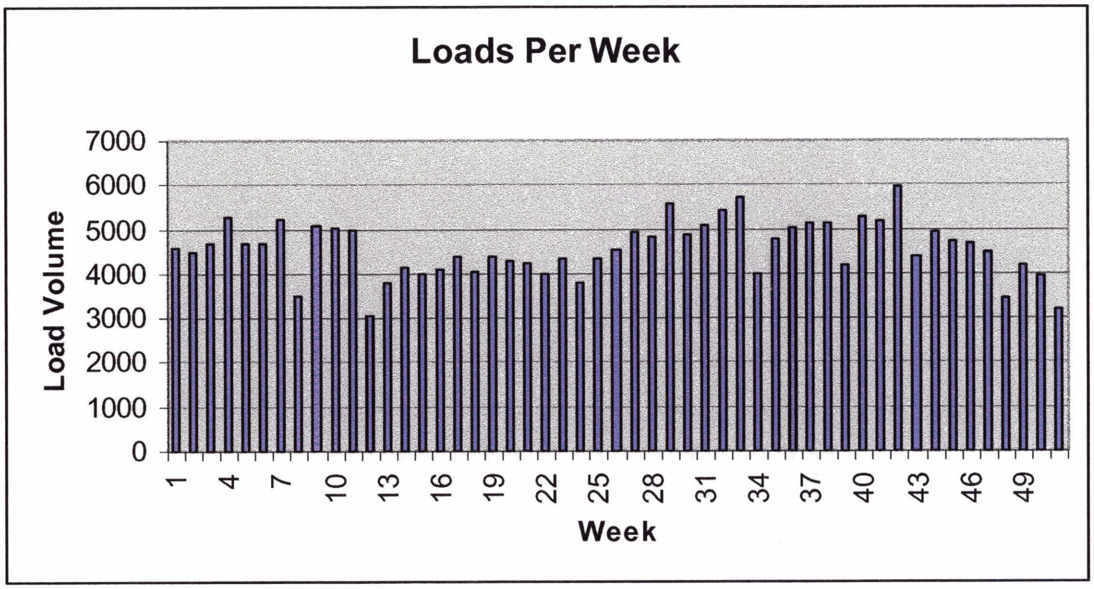

Figure 8 - Load Volumes by Week

2.3.2 – Load Volume by Lane

Average load volume per lane was 126.6 loads over the 51 week period, or 2.48 loads per week. Lane volumes varied significantly by lane as shown in Figure 10, which shows the number of lanes categorized by total loads over the period. Here we see that nearly 800 of the 1838 lanes had 10 or less loads in the period.

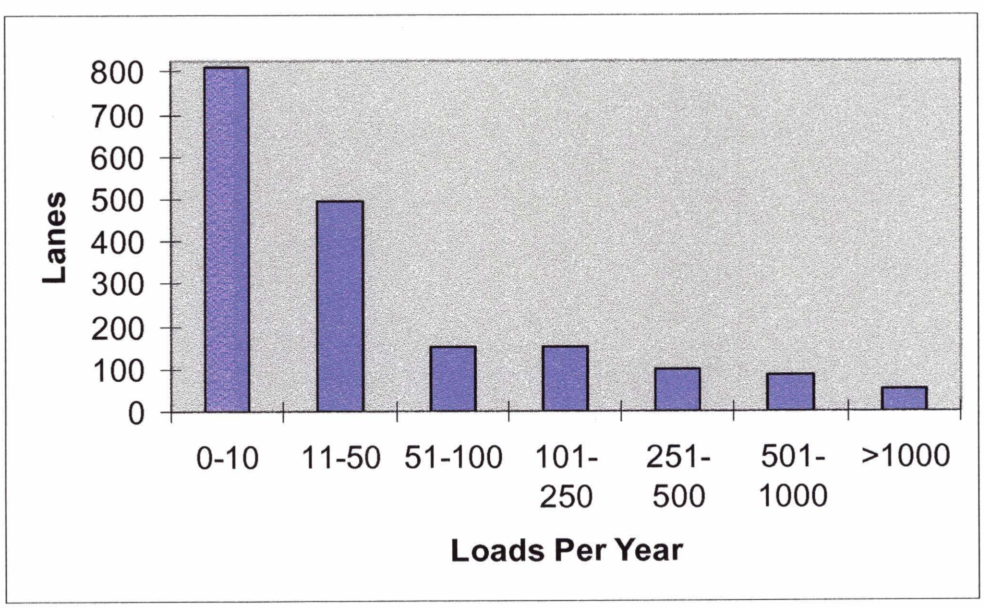

Figure 9 - Frequency of Lane Volumes

The statistical breakdown of the aggregated lanes shows that the mean volume per lane may be misleading if used inappropriately. The median number of loads per lane was just 14 with 44% of the lanes having 10 or fewer loads per year.

The high degree of variation in lane volumes indicates that there are a large number of lanes that account for relatively little in the way of total volume and relatively few lanes that make up a large percentage of the loads transported. In the data analyzed, 80% of the total load volume occurs on just 253 lanes, or just under 13.8% of the total lanes serviced by this shipper as seen in Figure 11.

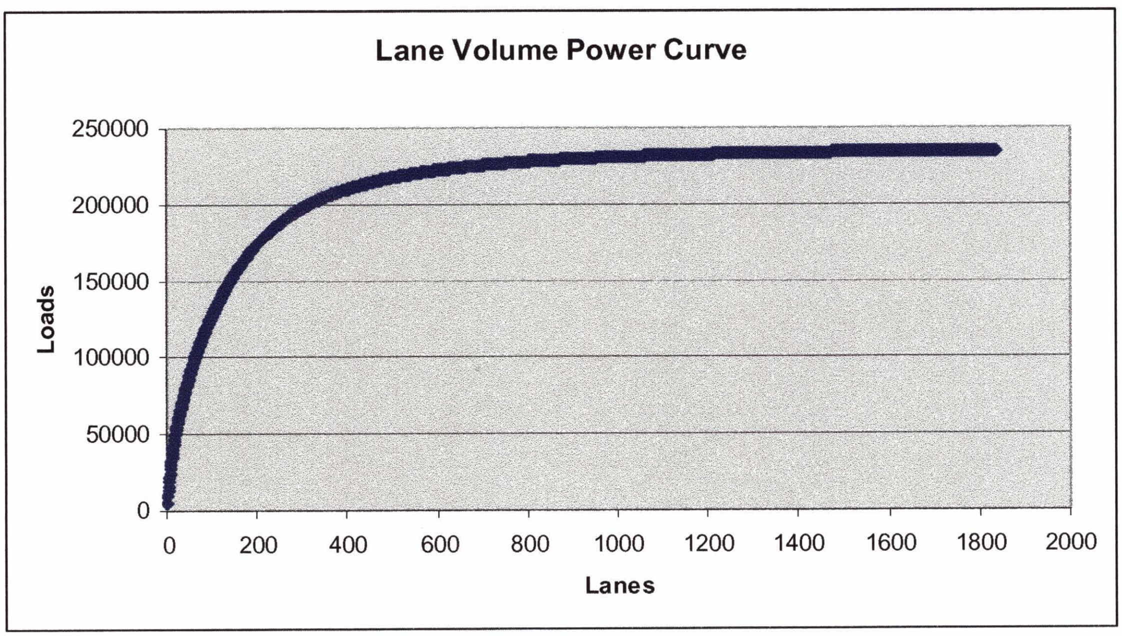

Figure 10 - Lane Volume Distribution

2.3.3 – Load Variation by Lane

When graphing the coefficient of variation against the total number of loads in the period, we can see a relationship. As lane volumes increase, the trend is for the variability to decrease. The correlation of lane volume to the coefficient of variation was

-.359, which indicates a moderate negative correlation between lane volume and the coefficient of variation.

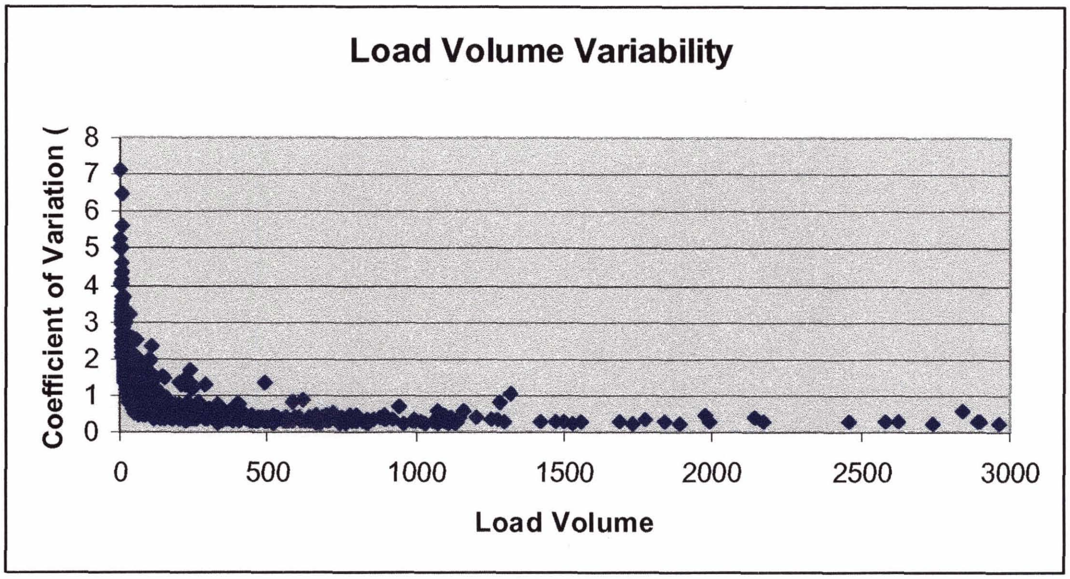

Figure 11 - Variability by Volume

Figure 12 shows the average coefficient of variation based on annual lane volumes. As expected, we can clearly see that the general trend that lanes with higher average volumes do tend to have lower variability. The lanes that have the highest coefficient of variance are those with annual load volumes of less than 100.

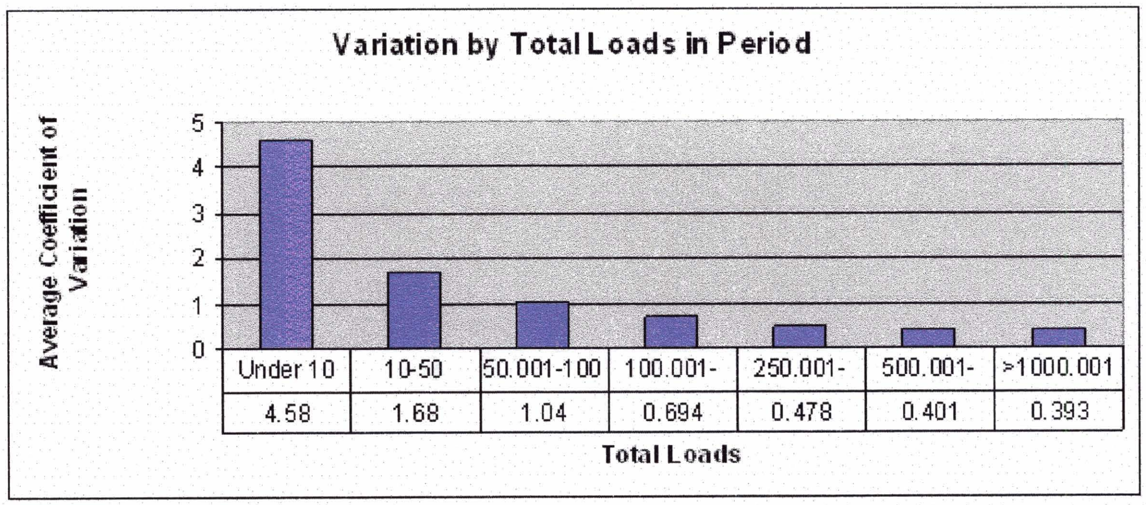

Figure 12 - Lane Variation by Annual Volume

2.3.4 - Variation by Load Distance

Private and dedicated fleets are often used on lanes that are less than 250 miles due to the fixed costs that are built into For-Hire carrier rates. Loads that are 250 miles or less also have the added benefit of enabling the driver to return to the domicile in one business day while adhering to the Hours of Service rules. A manager of a global 3PL said this when explaining how they make determinations in regard to which lanes should be utilized for their fleet.

> *"We usually start with a radius that depicts a one day round trip (250 miles or so) and then tailor the service area based on more populated delivery areas, then routes are built, service area is tailored again based on these routes, and then a utilization chart is built, then routes are tailored again to get optimal utilization. The above is a new business scenario. Operationally, it is usually determined which stores go private fleet and then only tailored once to add or remove compatible stops." (Mulqueen, 2006)*

Variability and how it correlates to the distance of loads was analyzed for the given data set. The loads that had the lowest overall variability were the loads that were under 250 miles.

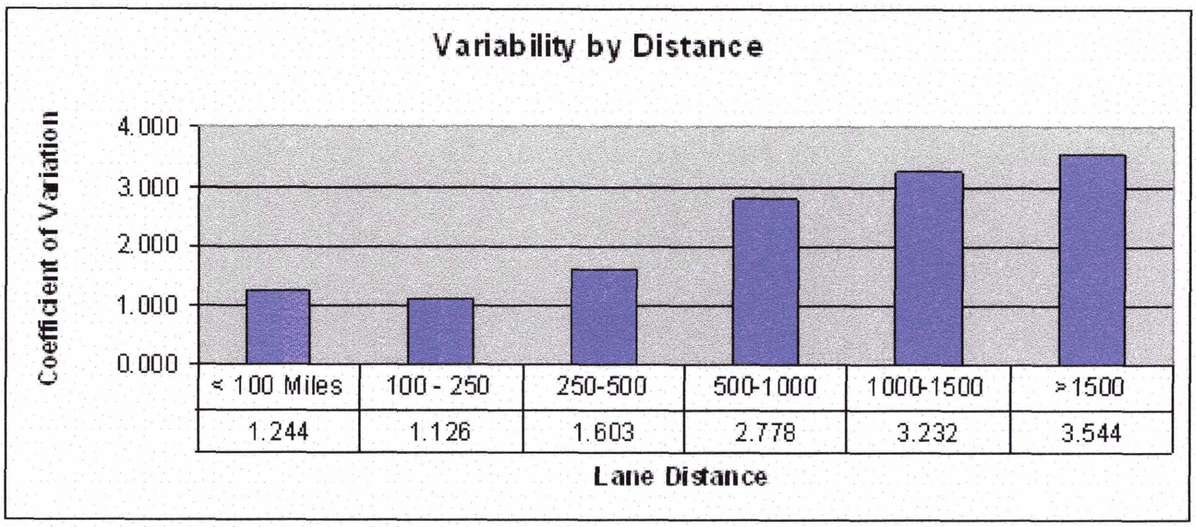

Figure 13 - Load Variability by Distance

2.3.5 Individual Lane Variability

While low volumes lanes are extremely volatile in terms of variability as defined by the coefficient of variance, even lanes with high volumes had a good deal of week to week variability. A high volume, but otherwise representative lane (12345) is shown below.

Lane - 12345	
Average	94.0
Median	86
Standard Deviation	29.5
Coefficient of Variation	0.31
Minimum	47
Maximum	174

Table 6 - Statistics for Lane 12345

Weekly fluctuations are relatively pronounced with volume spikes on multiple occasions generating volumes that exceed 100% of the volume of the previous week and then returning to more normal levels the subsequent week after the spike (figure 15).

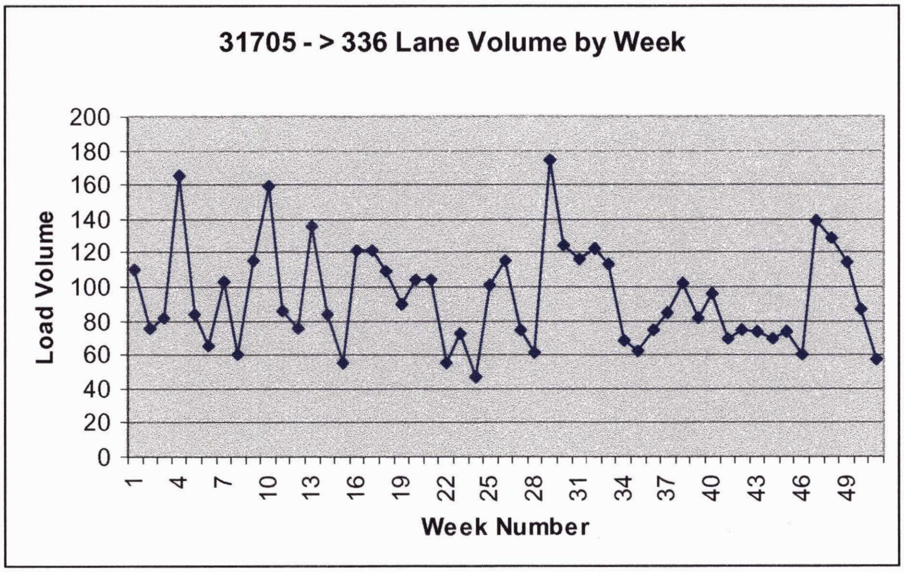

Figure 14 - Individual Lane Variation

2.4 Data Characterization Summary

The analysis of this CPG shipper data show that a large majority of the freight lanes serviced in the network are low volume and sporadic, so would not be good candidates for the fleet and should be removed from consideration during the tour formation and assignment process. Even attractive lanes in terms of load volumes and consistency have relatively high levels of weekly load volume variability. A transportation policy that does not take into account this variability in its construction will provide a sub-optimal solution for the shipper.

When analyzing the historical rates charged for the CPG shipper's loads, there is a high degree of cost/mile rate variability even for lanes that have similar attributes in terms of lane volume, origin point and distance covered. While no conclusive evidence exists that explains the rate variation, a comparison of rates for lanes that did and did not terminate in a known origin location backs up the premise that economies of scope, rather than economies of scale, are more important drivers when explaining rate variations.

3 – Methodology

3.1 – Overview

This chapter provides an overview of the methodology used to generate an enterprise's transportation policy. First, I will discuss the assumptions and methods that are used in the heuristics as well as within the process to formulate the policy.

I will then discuss and describe the steps that I am proposing to create the policy, providing details into the following processes:

- Lane aggregation and attribute definition
- Lane Segmentation
- Tour Identification
- Tour Asssignment/Simulation

3.2 Assumptions and Inferences

The framework is designed to help define the appropriate transportation policy for an enterprise. Due to the complexity of the problem, assumptions have been made and implemented within the optimizer. These assumptions are described below.

3.2.1 – Load Volume Pickup and Delivery Dates

This analysis assumes loads that occur within a weekly period are fluid in terms of load pickup and delivery days however all load volumes must be serviced within the specified period. For instance, if on a given lane, there are 10 loads per week, those loads can be picked up and delivered at any time within the week. More granular, load specific pickup and delivery windows are not supported by the model. The underlying inference from this assumption is that the transportation organization can drive the customer and/or buyer in terms of when product needs to be distributed.

3.2.2 – All Movements are Full Truckloads

The model assumes that all loads are full TL movements. Additionally, no intermediate stop-offs are supported in this model. The assumption is that there is one pickup and one delivery point per load.

3.2.3 – Pickup and Delivery Times

In determining the duration of a tour, time is not allocated for each pick-up and delivery. While this may seem counter intuitive, it was done as to not over penalize the fleet in terms of tour duration. The HOS rules allow a driver to work a total of 14 hours, but drive only 11. The 11 hours of drive time is represented in the heuristic by a maximum number of miles allowed per day. This number is specified at 500 miles per day throughout the analysis. The assumption is that the remaining time in a driver's day that is not included in the drive time will be used to load/unload the vehicle at the load origin and destination points.

3.2.4 – Accessorial Charges

Accessorial charges are incurred by the shipper when For Hire carriers are asked to perform services that are beyond the services included in the freight contract. Examples of accessorial charges include lumper fees, stop-off charges, demurrage and detention. The model does not take into account accessorials, which will have a tendency to bias the heuristic in favor of For Hire carriers since these costs are normally more pronounced if the services are done by a For-Hire carrier.

If a shipper does have large accessorial expenses, it would be recommended that when using this approach, they be estimated and included in the For Hire costing for each lane serviced in order to provide a better estimate of the true cost of using For Hire carriers.

3.2.5 – Equipment Compatibility

When forming a tour, the heuristic assumes that loads do not have unique vehicle requirements that would make the equipment used within a tour incompatible with the equipment required for any proceeding or subsequent leg on that same tour. For example, the model would not be able to handle a situation whereby leg 1 required a dry van and leg 2 required a flat bed. Additionally, if leg 1 required a 53' trailer and leg 2 required a 28' pup due to customer receiving limitations, the model will not provide the proper solution. The assumption used here is that the trailers used are uniform in functionality and size or that drop trailers are used, whereby the appropriate equipment is positioned to support the requirements of the subsequent leg in a tour.

3.3 Determining the Lane Distribution Type

In order to appropriately define the expected lane volumes given pre-specified user confidence levels, a determination needs to be made as to which type of distribution most closely represents the weekly lane volumes over the given period. To determine this, expected lane volumes are calculated based on assuming normal and poisson distributions. We then can use a Chi-Squared test to determine the probability that the null hypothesis – i.e. that the observed lane volumes approximate the tested distribution - is true.

For example, Lane ABC has the following attributes:

Average	0.35
Standard Deviation	0.52
Weeks	51.00

Based on these values, the anticipated number of weeks that a specified number of loads will appear can be specified.

Loads Per Week	*Normal Distribution Expected Value*	*Poisson Distribution Expected Value*
0	12.7326	35.8335
1	32.7718	12.6471
2	5.4544	2.2318
3	0.0412	0.2626
4	0.0000	0.0232
5	0.0000	0.0016

Table 7 - Expected Value Variation by Distribution Type

In this example, we can see that over a 51 week period the properties of a normal distribution specify that 12.7 weeks will have 0 loads and 32.7 weeks will have one load and so on. Similarly, using a Poisson distribution, 35.8 weeks will have 0 loads and 12.6 weeks will have 1.

In order to determine which distribution type more accurately portrays the data in question, we can compare the values actually observed in the data against each distribution type.

Loads Per Week	*Normal Distribution Expected Value*	*Poisson Distribution Expected Value*	*Actual Data*
0	12.7326	35.8335	34
1	32.7718	12.6471	16
2	5.4544	2.2318	1
3	0.0412	0.2626	0
4	0.0000	0.0232	0
5	0.0000	0.0016	0

Table 8 - Comparing Actual Data Against the expected values of each distribution type

By using Excel's chitest function, we can now determine that probability that the null hypothesis for each distribution type is true. In this example, the chitest returns the following values:

Normal Distribution	***2.365E-10***
Poisson Distrubition	***0.8560***

In this example, the Chi-Square test results show that the null hypothesis for a Normal distribution must be rejected since the returned value is extremely low, however, the test for the Poisson distribution tells us that there is an 85.6% probability that the null hypothesis is true as we compare the observed data with the expected data given a Poisson distribution.

Figure 16 charts the observed and expected values. We can see how much more closely the actual values align with a Poisson versus Normal distribution. .

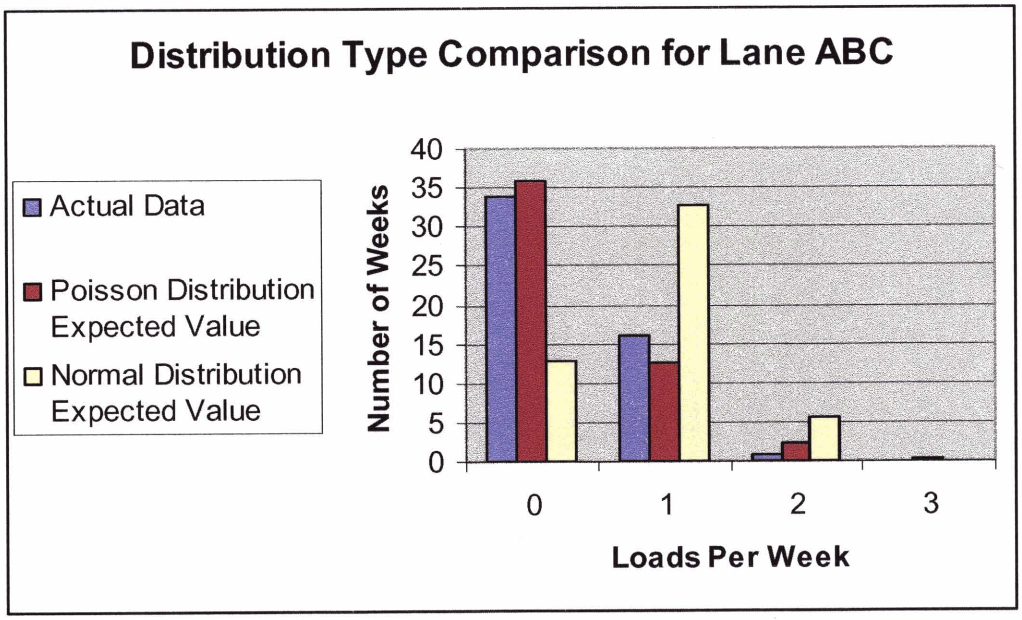

Figure 15 - Assigning the Proper Distribution Type

It is also important to understand the ramifications of a misidentified distribution type. For a load with a mean of 14, and standard deviation of 5 assuming a normal distribution, the curves are markedly different.

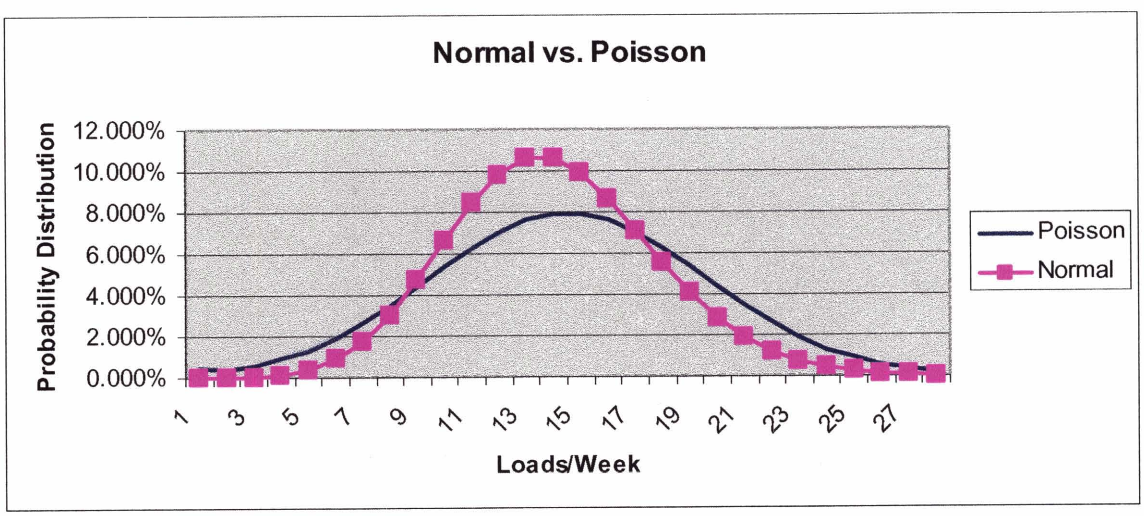

Figure 16 - Visualizing the Difference in Distribution Types

In the proposed process, a misidentified distribution type will cause incorrect values to be used during optimization and simulation, leading to an incorrect solution. Table 9 shows the actual values that would be used for various confidence levels, which gives an indication of the potential for incorrect solutions if the time is not taken to properly identify distribution types. Distributions that are identified as normal yet are actually poisson will be over-allocated load lane volumes up until quantities reach the the mean, at which time, they will be under-allocated load volumes.

Confidence	Loads within confidence - Normal	Loads within confidence - Poisson	Normal Value – Poisson Value
99%	5	1	4
90%	8	5	3
80%	9	6	3
70%	10	8	2
60%	11	10	1
50%	12	12	0
40%	13	14	-1
30%	14	16	-2

Table 9 - Poisson vs. Normal Confidence Values

3.4 – Network Aggregation

The network aggregation process is simply a way to group load volumes together in a meaningful way, thereby specifying the periodic load quantities for the lane in question. In the example of the consumer goods company discussed in the previous section, loads that originated at a common location, but terminated within the same 3 digit postal code were aggregated together. Aggregation logic based on similar geographic load characteristics could be made more specific by combining loads that are within the same 5 digit postal code or even same pickup and/or delivery address. However, the more granular the aggregation logic, the more lanes will be created, each with less load volume than those lanes with higher levels of aggregation. This will make the solution more complex without adding significant operational benefit.

Another possible way to aggregate would be to include a temporal component whereby only loads with overlapping pickup and delivery windows would be aggregated. Again, this vastly increases the number of lanes and lowers the numbers of loads/lane making tour formation reliant on existing delivery windows. Only networks that have extremely large lane volumes would support this level of aggregation during the tour formation process.

3.5 – Lane Segmentation

The role of the lane segmentation process is to make a determination as to which of the lanes in the distribution network should be considered as potential lanes in the tour formation process. By pruning off lanes that do not have the proper characteristics for tour formation, we are able to reduce the number of lane options that the tour generation process must consider. The segmentation process requires the use of a quantitative approach whereby various lane attributes are weighed. To do this, we must come to agreement as to the attributes of a lane that makes it worthy of tour formation consideration.

From a business perspective, lanes that are most attractive to a private fleet are ones that have the following attributes:

- The lane has a high volume of loads over the period of time the transportation policy will be in effect
- The lane has a high level of consistency over the period

- ➢ The relative cost of using For Hire Carriers vs. the Private Fleet is high

Additionally, the business may require the use of private/dedicated fleet on various lanes for qualitative reasons, such as increased levels of service and/or operational control.

3.5.1 – Segmentation Inputs

Table 10 specifies an example of a data record used in the lane segmentation process. In this example, every week under consideration for a given lane will have the following attributes, although not the same data.

Lane Attributes	*Example*
Origin	Boston, MA
Destination	Salt Lake City, UT
Time Period	Week 1
Load Volume	30
Contract Cost / Load	$1750
Fleet Cost / Load	$1400

Table 10 - Lane Attributes

This lane will be grouped with records from other time periods that share the same origin and destination. From this grouping, we will be able to determine for each lane:

1. The lane's mean weekly volume
2. The lane's weekly standard deviation
3. The lanes distribution type (Normal, Poisson, Other)
4. The lane's Coefficient of Variation
5. The ratio of Contract Cost over Fleet Cost

The segmentation logic will also utilize input data that specifies the minimum lane volume within a confidence level. This will enable the segmentation logic to filter out lanes that do not have the necessary volume or reliability to support a user defined minimum number of loads per lane

For example, to consider a lane for tour formations, a planner may decide that the policy used needs to ensure that there is at least an 80% likelihood that the number of loads on a lane will equal or exceed 5 loads per week. Table 11 shows how this would be resolved in the segmentation process.

Lane	***Mean Weekly Lane Volume***	***Standard Deviation***	***Distribution Type***	***Number of Loads at 80% Confidence***	***Consider for Tour Formation***
Lane A	12	3	Normal	9.48	Yes
Lane B	8	1	Normal	7.16	Yes
Lane C	9	3	Poisson	6.00	Yes
Lane D	15	12	Normal	4.90	No
Lane E	9	6	Normal	3.95	No

Table 11 - Segmentation Logic

Because the planner indicated that a minimum of 5 loads per week are required, only lanes A, B and C would be considered for tour formation. Interestingly, in this mock example, it should be noted that the lane with the highest mean weekly volume did not meet the planner's criteria. This is an example of how utilizing an average without understanding the variability associated with that average can return misleading and sub-optimal results.

To better illustrate variability within a lane, let us look at a lane in more depth. Assume the following attributes represent the lane's historical load volumes.

Lane Attribute	Value
Weekly Mean	32
Standard Deviation	4
Distribution Type	Normal
Coefficient of Variation	.125
Contract Cost / Fleet Cost	1.25
Confidence Level	.90

Table 12 - Lane Attributes

Using this data, we are able to determine the probability that within a specified confidence level, the minimum required weekly lane volume will be met. Let us assume that the minimum weekly lane volume specified by the planner requires that with a 90% confidence level, there will be at least 25 loads available per week.

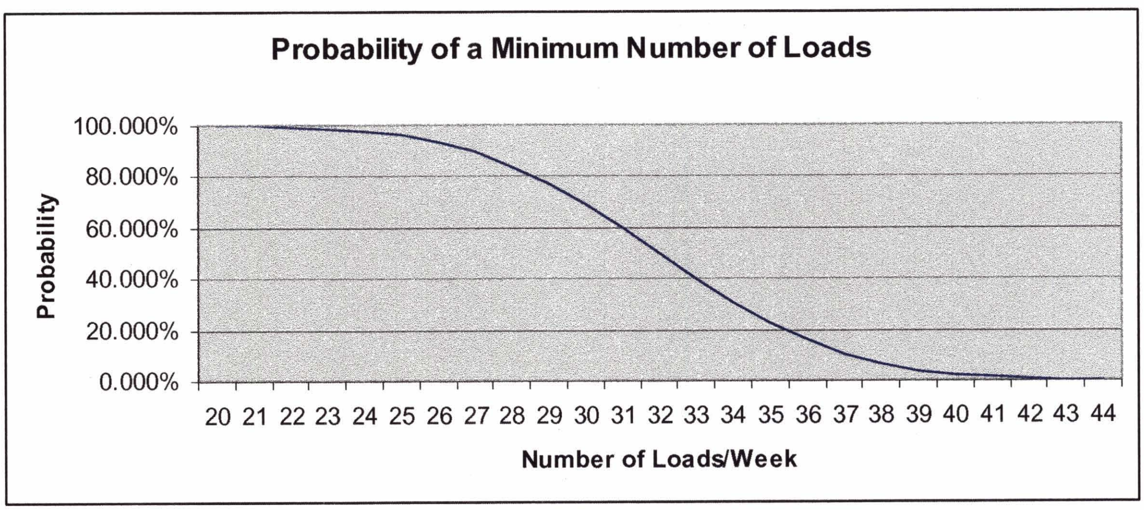

Figure 17 - Probability of X number of loads

By calculating the inverse of the normal distribution we are able to compute that the probability of at least 25 loads per week occurring equals 95.99%. Figure 19 shows another view of the same data. The area in black represents the probability that the weekly volume will be less than 25.

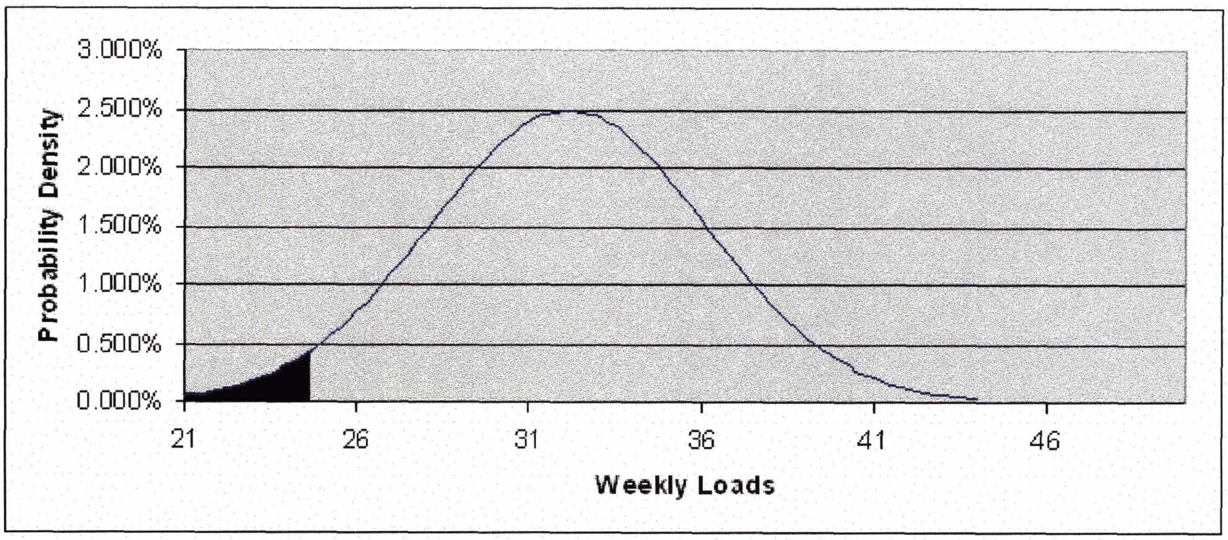

Figure 18 - Normal Distribution Curve that indicates 95.99% Confidence

3.6 – Tour Identification

This thesis will not provide insight into heuristics or methodologies that can be used to identify tours in large networks, but will assume that a viable method to identify potential tours that the optimization routine should consider when assigning load volumes will be applied. Tour formation heuristics has been addressed, by among others, Arunapuram (1993) and Caplice (1996).

The assumption in this work is that there are a limited number of tours that can be pre-identified along with the expected cost savings that each single tour volume will generate when compared against For-Hire costs. In order to limit the number of tours under consideration, as well as to ensure that only operationally feasible tours are considered, practical constraints are normally applied. Example of constraints, and the ones used in the identification of tours in Section 5 are:

- A Tour must originate at an Origin Node. An origin node is a location within the network from which load volume originates.
- A Tour must terminate at the same origin node as to which it started. This will generate a closed loop tour. This is valid in a private fleet environment, since it is incumbent upon the fleet operator to return drivers to their respective domiciles. It is not necessarily required if the load volume is going to be given to For-Hire carriers, since an open tour still may be valuable to the carrier.
- A tour may not revisit the same node more than 1 time in a single tour except when the tour is returning to the origin node, thereby ending the tour.
- A tour may not have 2 consecutive empty legs.

3.7 - Tour Assignment

Once the appropriate set of tours has been defined, the tour assignment logic's objective is to maximize savings by reducing the initial solution cost. The initial solution cost is established by assuming all lanes and their associated weekly volumes are serviced by For-Hire carriers. By assigning lane volumes to tours executed by the fleet, savings can be achieved if the total cost of the tour along with the fixed cost of the resource is less than the cost of servicing the individual lane volumes with For-Hire carriers.

As discussed in Section 2, private and dedicated fleets incur costs in fundamentally different ways than shippers incur costs when employing For-Hire carriers. When considering the fleet for a tour, the optimization needs to take into account the fixed fleet cost that is incurred regardless of use as well as empty miles, which are required to reposition the equipment from the destination point of a load to the origination point of the subsequent load, or the termination point of the tour.

In the tour assignment optimization routine, a fixed cost per period is defined, and tours are assigned a pro-rated portion of that cost based on the duration of the tour. For instance, if a tour's duration is 1/2 of a week, and the weekly fixed cost for a fleet resource is $500, the tour will incur a $250 fee for utilization of the equipment. This cost represents items such as fixed lease costs, depreciation, fixed labor costs and insurance. The per mile cost is calculated by multiplying the tour distance and a per mile rate. This cost is meant to represent items such as diesel fuel, tires, oil, maintenance and the variable portion of a driver's pay.

A final component of the optimization looks to maximize savings by utilizing the private fleet to generate revenue by moving third party freight on empty lanes. The optimization routine will increase total savings by multiplying a user specified value that indicates the probability of finding a third party movement on a given lane with the market rate of the given lane.

Figure 20 shows a four node network. The diamonds A and C are load origin points while rectangles B and D are load destinations. Using this example, we can show how the optimization routine would make its determination as to whether a fleet or For-Hire carrier should be used.

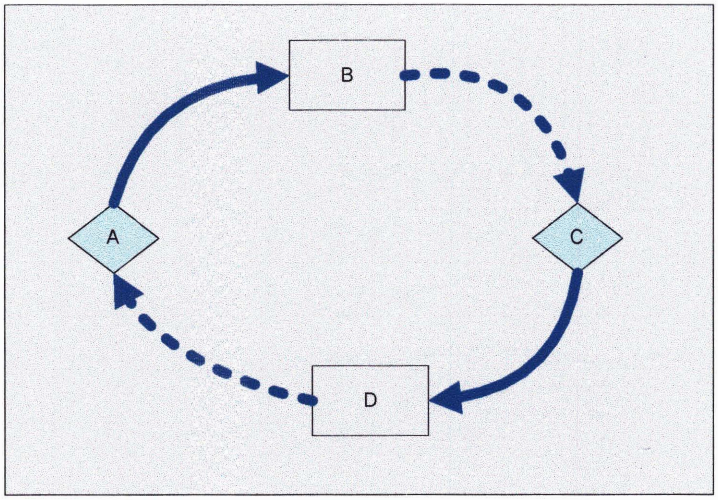

Figure 19 - Four node network

Assume the following:

- One load originates out of A destined for B
- One load originates out of C destined for D
- There is no volume between points BC and DA
- Each leg takes 1 day to complete
- All lanes are 500 miles
- The market rates for Loads AB and CD are $1000 each
- The fixed daily cost for a fleet resource is $100/day
- The cost/mile for the fleet is 1.00 for all loaded and empty miles
- The probability of the fleet getting a third party backhaul on lanes BD and DA is 0

In this example the For-Hire cost is equal to $2000. The fleet cost equals:

Distance cost ($1.00*2000 miles)

\+ Fixed cost (4 days * 100/day)

$2400

Since the fleet ($2400) is more expensive than the For-Hire carrier ($2000), the lane volumes will be given to the For-Hire carrier unless the lane volumes can be made part of another tour that does generate savings.

We created an Integer Program that finds the optimal number of tours that maximizes savings. The objective function of the optimization function is described below:

$$Max\ z = (\sum_{l \in L} FHC_l) - \left[(\sum_{k \in K} TR_k - TC_k) X_k \right] - FTC * Y + \sum_{l} \sum_{k} [(SLP_l)(X_k)(a_{lk})(C_l)] \quad (1)$$

subject to:

$$\sum_{k \in K} a_{lk} X_k \le Vol_l \quad \forall\ l \text{ that are loaded} \quad (2)$$

$$\sum_{k \in K} TD_k X_k - Y \le 0 \quad (3)$$

$$Y \le \text{Max Truck} \quad (4)$$

$$Y \text{ is an integer} \ge 0 \quad (5)$$

$$X_k \text{ is an integer} \ge 0\ \forall\ k \quad (6)$$

Where

X_k	= the number of trucks assigned to tour k
Y	= # of trucks in the fleet per week
FHC_l	= For Hire cost for lane l at the forecasted volume
FTC	= Fixed fleet truck cost per week
TC_k	=Fleet Truck Cost for Tour k
TD_k	= Tour Duration for Tour k
TR_k	= For Hire rate for Tour k
SLP_l	= Probability of a spot load on Lane l
L	= *Set of* All Lanes
a_{lk}	= 1 if $l \in k$, 0 if $l \notin k$
C_l	= For Hire Cost of Lane l
K	= Set of all tours

3.7.1 – Simulation

The tour assignment logic described above is provided as an input with a maximum lane volume based on user defined confidence levels. The lower the confidence level, the more loads are available to the tour assignment optimization routine. This approach completely disregards lane variability once the initial confidence factor is selected. In order to ensure that variability is taken into account when formulating the policy, modeling is used to assign simulated volumes to lanes over a 1 year period based on the historical lane volume and variability. This allows us to utilize the results of the optimization to see how well they stand up to the anticipated variability that exists within the network. By doing this, we are able to compare the simulated solution cost versus the optimal solution cost, which effectively provides us the cost of variability within the network. More importantly, it enables us to define how best to utilize fleet resources given the inherent variability of the network

3.8 – Plan Output

After optimization and simulation, the transportation policy is partially defined. At this point, the following decisions regarding the creation of transportation policy have been made:

- The size of the internal fleet.
- The lanes in which the fleet should operate.
- The lane volumes that the fleet will cover.
- The tours that will executed by the fleet.
- The anticipated load volumes that will require For-Hire carriers

4.0 – Data Analysis

This section of the thesis will step through the process discussed in the previous chapter using an artificial network created to illustrate the key insights this research has focused on. The most important of these, and the areas that the remainder of this portion of the research will focus, are the following

- The cost of variability and the danger of only using the average load volumes in networks with even low to moderate variability
- A discussion on finding the optimal confidence level to ensure maximum fleet utilization
- The power of even moderate usage of third party backhauls to generate revenue and take additional control of freight
- The potential of risk pooling fleet resources to maximize vehicle utilization

4.1 – Describing the Network

A simple, 5 node network is shown in figure 21. This network will be the basis for the analysis and insights that follow.

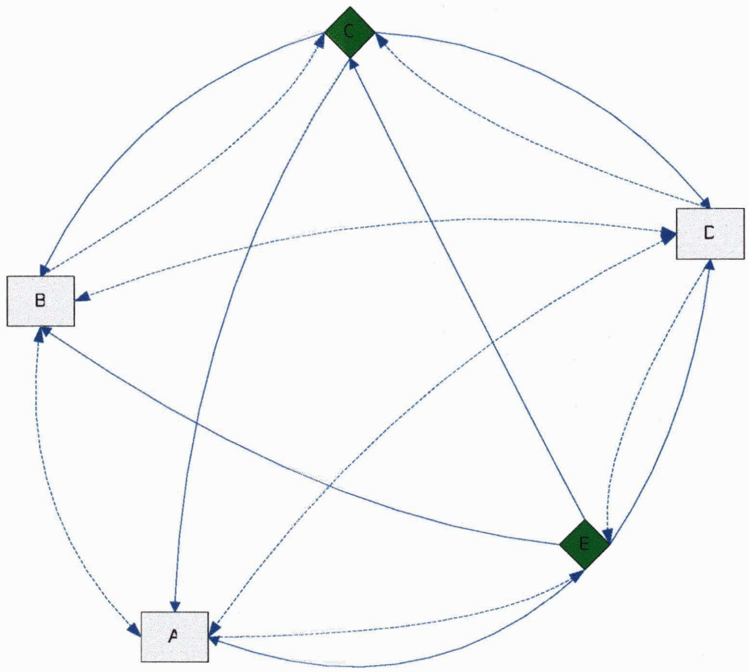

Figure 20 - Simple 5 Node Network

The solid arcs represent loaded lanes while the dashed arcs represent empty lanes. The diamonds (C,E) are the sole points of origin for the shipper's volume. The squares (A, B D) represent points that are strictly load destinations. However, as we will see later, through the utilization of the spot market, those points may be origin loads for third party movements.

In this network, there are 14 lanes. Key attributes of these lanes are described in table 13.

Lane	Type	Distance	Market Rate	Weekly Volume	Market Cost/Load	Trip Duration (Days)	Fixed Cost	Variable Cost
A to C	Empty	1000	$ 2.00	0	$0	2	$400	$800
A to E	Empty	600	$ 2.25	0	$0	1.2	$240	$480
B to C	Empty	500	$ 2.40	0	$0	1	$200	$400
B to E	Empty	1100	$ 1.60	0	$0	2.2	$440	$880
C to A	Loaded	1000	$ 1.80	12	$1800	2	$400	$800
C to B	Loaded	500	$ 2.10	14	$1050	1	$200	$400
C to D	Loaded	600	$ 2.20	13	$1320	1.2	$240	$480
C to E	Loaded	800	$ 1.65	21	$1320	1.6	$320	$640
D to C	Empty	600	$ 2.30	0	$0	1.2	$240	$480
D to E	Empty	400	$ 1.90	0	$0	0.8	$160	$320
E to A	Loaded	600	$ 2.50	16	$1500	1.2	$240	$480
E to B	Loaded	1100	$ 1.75	11	$1925	2.2	$440	$880
E to C	Loaded	900	$ 1.90	15	$1710	1.8	$360	$720
E to D	Loaded	400	$ 2.70	12	$1080	0.8	$160	$320

Table 13 - Lane Attributes

Data Descriptions

1. Lane – Defines the origin and destination of a movement
2. Type – Defines whether the leg is loaded or unloaded. For Hire carriers are not costed for traversing empty legs.
3. Distance – The length of the lane in miles
4. Market Rate – The standard rate shippers get charged on the lane. These rates do not apply for empty lanes, except when calculating backhaul revenue for tours
5. Weekly Volume – The mean weekly volume for the lane. Empty lanes have no volume, since the shipper does not have loads that originate from those points.
6. Market Cost / Load – The market rate multipled by the distance. Only applies to For-Hire carriers

7. Trip Duration Days – The period of time required to service the lane. This is calculated by dividing the total lane distance by 500 miles, which is the daily number of miles per day the fleet is expected to travel.
8. Fixed Cost – The fixed weekly charge, multiplied by the duration of the lane.
9. Variable Cost – The length of the lane multiplied by the fleet's per mile cost.

4.2 – Tour Identification

An input into the optimization is the specification of all valid tours that are to be considered when assigning load volumes. In this example, a Depth-First Search (DFS) is used to identify tours. To reduce the number of potential tours, as well as to ensure that only operationally feasible tours will be considered, the following logic is applied.

1. Tours must originate at an origin point (C or E)
2. Tours must terminate at the same point in which they started
3. Tours are not permitted to have 2 consecutive empty legs
4. Tours cannot return to the same point more than once, except when returning to the origin point to terminate a tour

The DFS resulted in the following 32 potential tours as shown in figure 22.

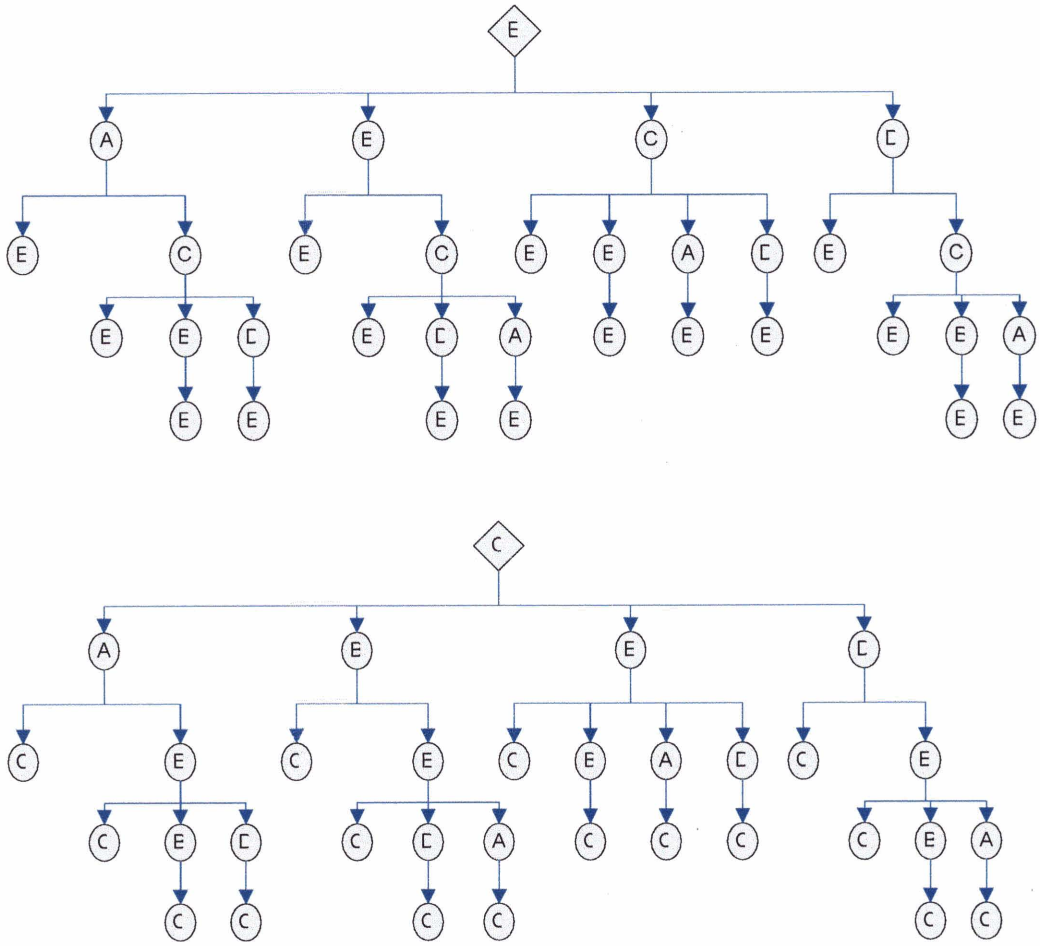

Figure 21 - Depth First Search Visualization

By implementing this logic, the number of tours the optimizer will need to consider is limited to 32, with 16 tours originating from both point C and point E. Once the valid tours are created, we have the ability to compare the anticipated tour cost against the For-Hire cost to determine a net savings achieved when using a single fleet resource on any given tour. It is this value that is used within the objective function of the optimization routine to determine the volume on each tour. Based on the values given in table 14 and the costing methodology described in section 4.5, the expected savings for the 32 tours in question is as follows:

C Origin Tours			E Origin Tours	
Tour	Savings		Tour	Savings
CAC	-$600		EAE	$60
CAEC	$510		EACE	-$60
CAEBC	-$115		EACBE	-$1,290
CAEDC	-$240		EACDE	-$300
CBC	-$150		EBE	-$715
CBEC	-$240		EBCE	$365
CBEDC	-$990		EBCDE	$125
CBEAC	-$1,290		EBCAE	-$115
CEC	$990		ECE	$990
CEBC	$365		ECBE	-$240
CEAC	-$60		ECAE	$510
CEDC	$240		ECDC	$510
CDC	-$120		EDE	$120
CDEC	$750		EDCE	$240
CDEBC	$125		EDCBE	-$990
CDEAC	-$300		EDCAE	-$240

Table 14 - Tour Savings

Negative values indicate that the fleet costs are greater than the costs that would be incurred if a For-Hire carrier was used. This occurs when there is significant under utilization of the fleet equipment in terms of excess empty miles and/or when the For-Hire carrier has provided unusually low rates on lanes. Positive values indicate that utilization of the private fleet on that tour saves money versus the use of For-Hire carriers.

4.3 Specifying the Confidence Level

The confidence level is used to specify the lane volume that the optimizer uses during tour assignment. Based on the distribution type that was identified for the lane by the Chi-Squared test, we can calculate the inverse of the cumulative distribution and determine, based on the confidence level defined, the weekly lane volume that will provide the specified level of assurance.

For example, assume a lane is normally distributed and has a mean of 20 and a standard deviation of 10. If a confidence level of 50% is assigned, 20 loads would be made available to the optimizer for that lane. The 50% confidence level indicates that there is an equal probability that the number of loads per week will be greater than or less than 20.

4.4 - Simulation

To show the impact of variation, the transportation policy created by the optimization routine is passed to a simulator. Unlike the optimizer, which operates on a single, discrete value, the simulator generates representative lane volumes designed to emulate a period of time based on the historical lane volumes and distribution patterns. The goal of the simulation is to see how well the optimized transportation policy stands up to the variability that will inevitably occur within a network. For this process, the simulation was designed based on the following assumptions and rules:

1. Fifty-two weeks are simulated
2. Lane volumes fluctuate based on the each lane's specific distribution type, mean and standard deviation
3. If the required volume to complete a tour is not available, all other loaded legs on the tour assume For-Hire carriers executed the load(s). The fleet capacity originally allocated to that tour is left idle. While the fleet will not incur the variable costs associated with this tour volume, it will still

incur fixed costs. The assumption is that one broken leg on a tour nullifies the entire tour.

4. The simulation assignment logic will place all lane volume on the current tour before assigning any lane volume on the next tour under consideration. The tour assignment sequence is fixed and biased toward assigning the tours that generate the most savings first.
5. The fleet cost is fixed based on the number of vehicles required to handle the tour volume recommended by the optimizer
6. The fleet variable costs and For-Hire costs are calculated by running the simulation routine 10 times. The average of these costs is then taken. This is required because as lane volumes fluctuate, the For-Hire costs change as well as the distance based costs associated with the fleet.

4.5 – Optimization and Simulation Results

Analysis was done to study the potential impact that four specific areas have on the creation of transportation policy. These areas are:

1. The use of confidence levels to define lane volumes available for optimization
2. The use of simulation to validate the results returned from deterministic optimization solutions
3. The use of risk pooling to increase the utilization of fleet equipment
4. The use of third party backhauls to increase fleet utilization

4.5.1 – Varying Confidence Levels

Based on the values presented above, optimization was run using increasingly less restrictive confidence levels. Additionally, the inputs into the optimization utilized a Coefficient of Variation, as defined by the $\frac{\sigma}{\mu}$, of .25 for this analysis. The optimized results are shown in figure 23.

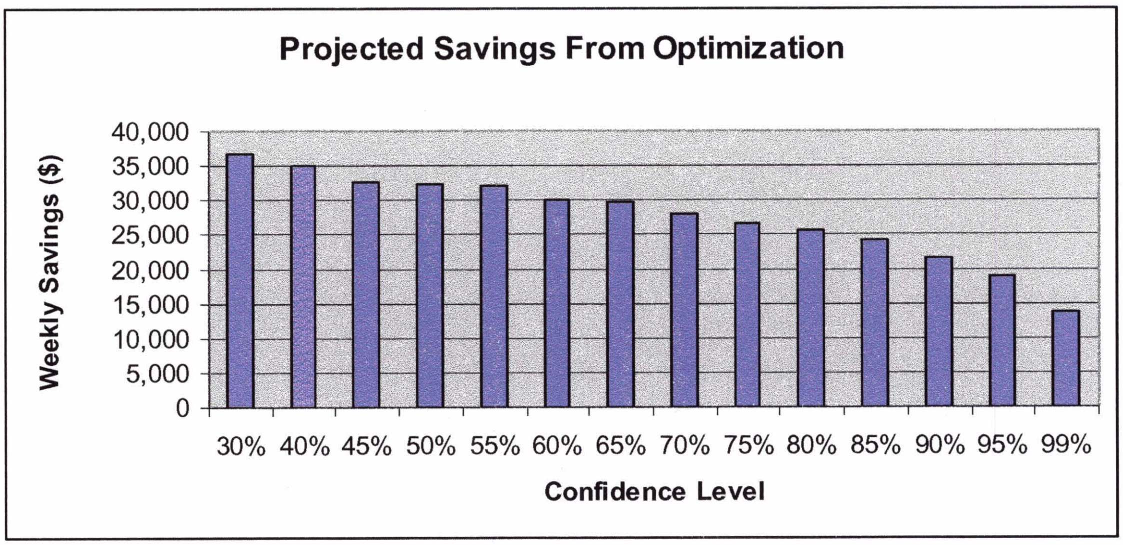

Figure 22 - Confidence Levels and Savings

As would be expected, as confidence levels increase, the projected weekly savings decrease. This is due to the decrease in lane volumes that are made available to the optimizer for tour assignment as the restrictions imposed by the confidence levels are tightened. Being deterministic, however, the optimizer does not take into account the cost of lane variability that will have an impact on the projected savings. This will be discussed in more depth later in this chapter.

In this example, the simulation shows a very different result than that of the optimization.

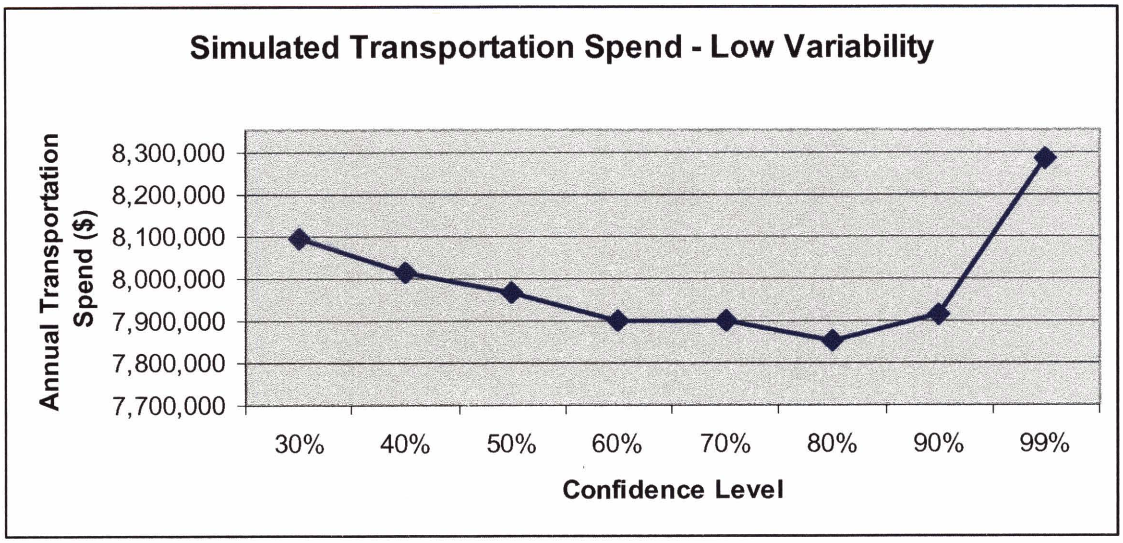

Figure 23 - Transportation Spend as it relates to the confidence level

When utilizing a very high confidence level, the overall solution costs remain high. This is due to the low lane volumes that are necessary to guarantee high confidence levels, thereby providing the optimization routine little chance to create tours and generate savings. As the confidence level is loosened, more tours are formed, and the overall solution cost is reduced. However, once the confidence level goes below 80%, the total cost of the solution begins to rise. This is due to an increase in the number of tours that were anticipated to form by the optimizer, but which did not have sufficient volumes generated during simulation.

4.5.2 - The Effects of Lane Variability

Lane volume variability plays a significant impact on the ability of the optimizer to generate savings in the tour assignment process. A high degree of variability reduces the number of lanes available for tour formation when applying confidence factors. Essentially, the distribution curve flattens out as the coefficient of variation increases, thereby reducing the volume within the specified confidence level. This is shown in figure 25.

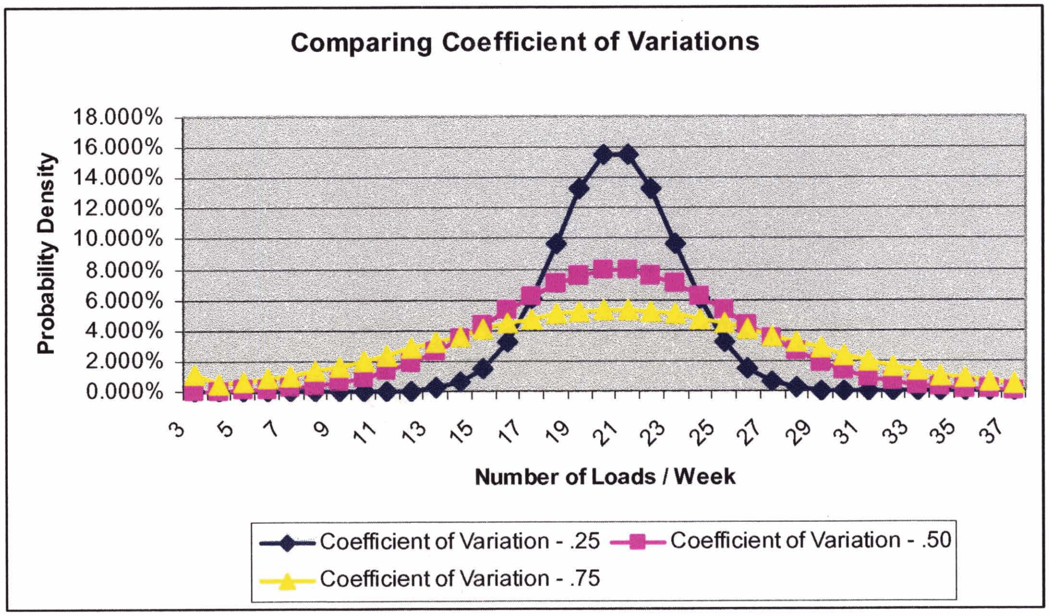

Figure 24 - Distribution Curves varied by C.V.

Because of this effect, optimization will not be able to generate as much in the way of potential savings as would be possible if more loads were made available. This was seen when contrasting the anticipated savings calculated during optimization and simulation when the coefficient of variation was increased.

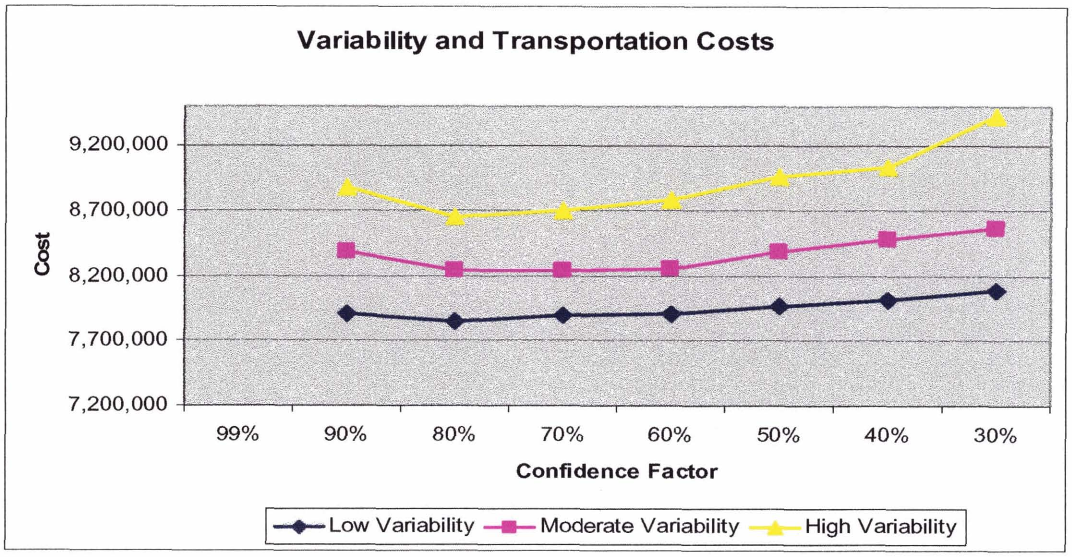

Figure 25 - Affects of Variation on Savings

As expected, the larger number of loads available for tour assignment when the variation was smaller enabled the optimizer to generate higher savings, thereby reducing the total transportation cost. As the coefficient of variation was increased, the total transportation cost increased.

In each case however, there was a minimum annual transportation cost that was reached before the total solution cost began to rise due to under utilized fleet resources. In each case, the minimum was reached when the confidence level was at or near 80%. It should not be inferred that the 80% value seen in this example in any way is a universal example, but is data specific and related to the ratio of fleet and For-Hire costs.

4.5.3 – Deterministic vs. Probalistic Solutions

Transportation decision support systems that utilize solely average based, deterministic methods at arriving at solutions whereby variability is prevalent should be viewed cautiously. Lane volume variability has a profound affect on costs and if ignored, will lead to a transportation policy that is only optimal when the forecasted lane volumes are 100% accurate, but sub-optimal when variability is taken into account.

Figure 27 overlays the simulated and optimized solutions showing the potentially misleading results of a deterministic optimization solution that does not take into account the variability within a distribution network.

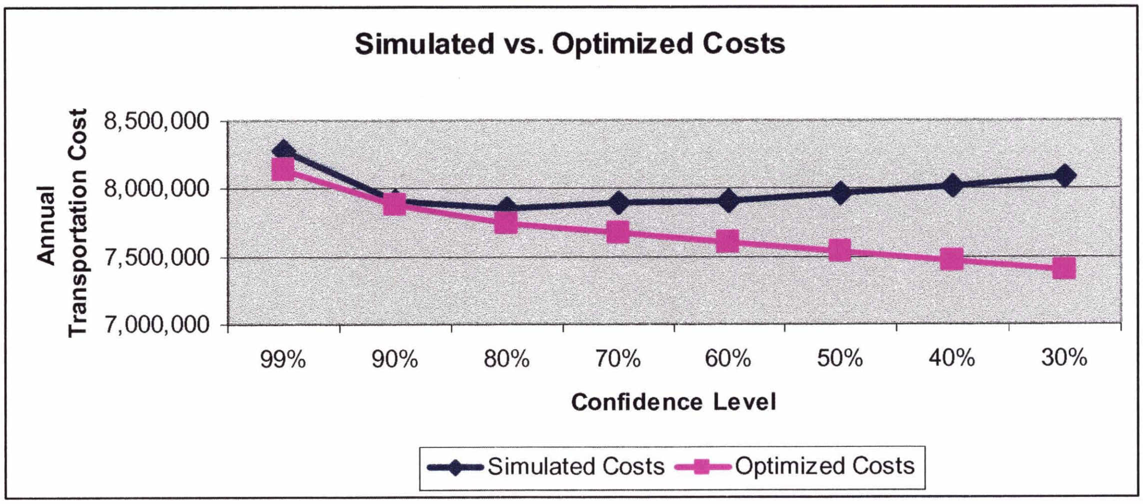

Figure 26 - Differences between Simulated and Optimized Costs

When confidence levels are high the results are relatively close. However, as confidence levels decrease, the anticipated savings of the optimization vis-à-vis simulation begin to diverge. The optimized solution anticipates ever increasing savings as the confidence levels decrease, while the simulated solution reaches a minimum cost, and than begins to trend upwards as costs increase due to a drop in fleet utilization. It should be noted that when using the average volume on a lane is equivalent to a 50% confidence level. In this example, the use of the average lane volume with regard to the mean yields a result that is 5.39% lower than the simulated findings.

4.5.4 - Tour Reliability

As equipment utilization falls below a certain threshold, the total solution cost begins to rise. This can be seen in Figure 28, which overlays the tour reliability, defined as the percentage of tours that are executed versus tours that the optimizer planned, but were never executed due to the anticipated volumes not appearing during the simulation. Here, as in the other analysis, the initial savings generated by using fleet resources on the most profitable tours outweighs the cost of running the fleet. This is because the fleet resources will be allocated to the lanes that generate the most revenue, and with still relatively high confidence levels, will most likely to executed. As the confidence level is decreased and more load volumes are made available to the tour assignment process, tours that generate less savings must be generated while a higher likelihood of tours not generating at all is incurred. This is the basis for the inflection point that we see in each of these curves.

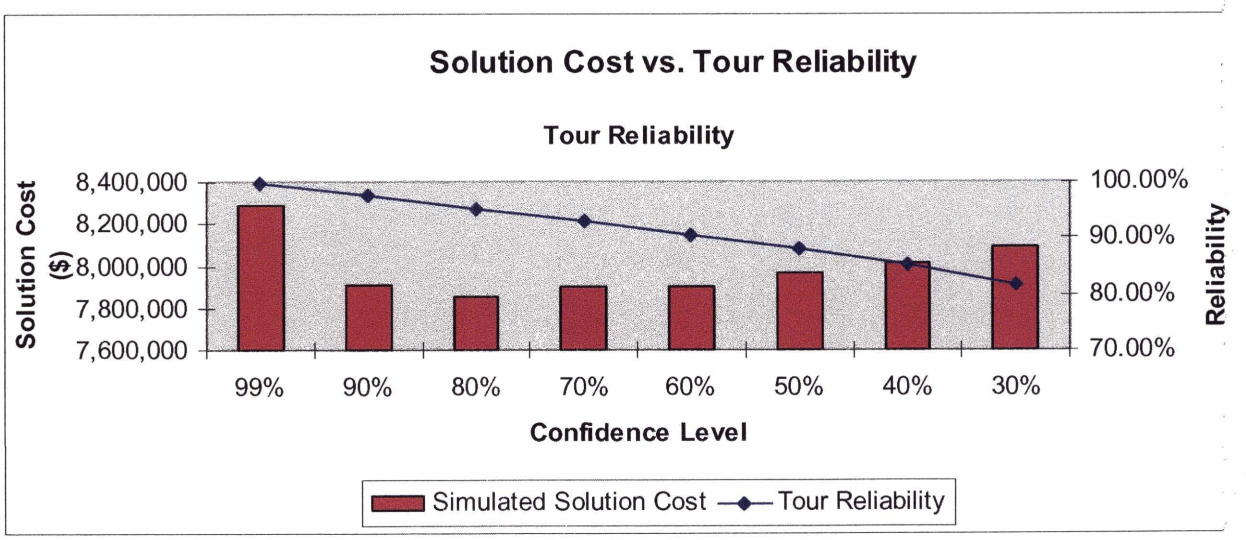

Figure 27 - Transportation Cost and Tour Reliability

In is interesting to note that even as the confidence level used drops to 30%, in this scenario, the reliability of the tours remains above 80%. Being somewhat counter intuitive, the reasoning for this phenomenon deserves an explanation.

At first blush it seems that if the lane confidence level is 30%, the reliability of a tour would not exceed 30% and that if in fact the lane volumes were independent, the probability of a tour would simply be the product of all the lane confidences.

However, as discussed in section 4.3.1, the confidence level is not the probability that any individual lane volume will occur, but the probability that the "last" lane volume will occur. If a lane's volume is normally distributed with a mean of 10 and a standard deviation is 5, there is a very high probability that the first, second and third loads will appear (96.4%, 94.5% and 91.9% respectively). It is as values get closer to the mean that the probability of each subsequent load not occurring becomes more substantial.

This can be visualized by looking at the total area under a normal distribution curve. At each of the tails, the incremental increase in area under the curve grows slowly as you move toward the mean. It is only as you are relatively close to the mean that the probability distribution density is large enough to have a real impact on decision making. It is within this area that you will typically find the tipping point in decision making, in this case, this point defines the optimal size of the fleet when compared against the use of For-Hire carriers.

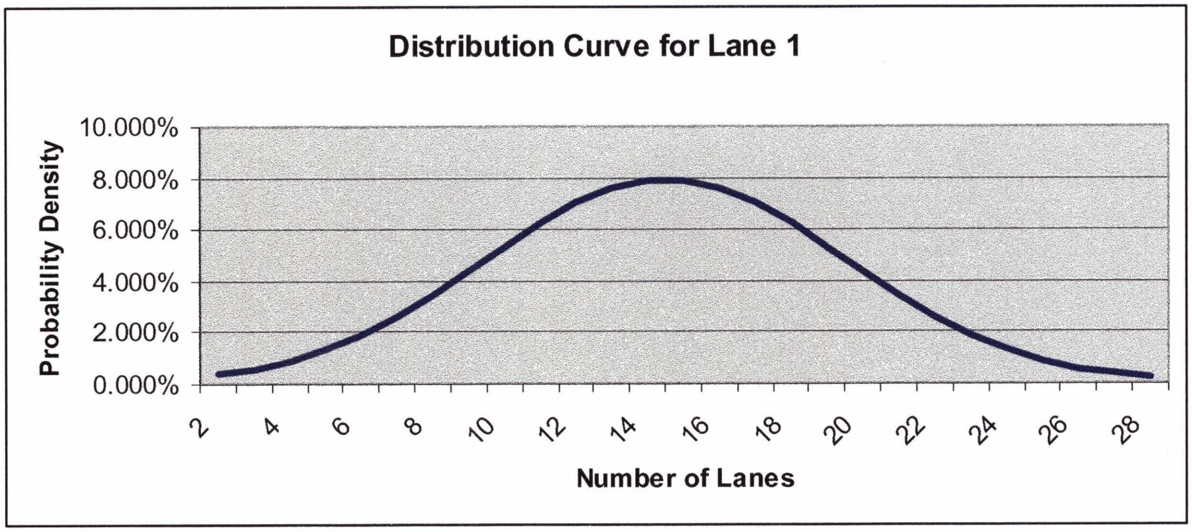

Figure 28 - Normal Distribution Curve

To see an example of this, consider the following scenario. Here we have two lanes. Based on the information provided, we can calculate the cumulative distribution of each of the lanes.

Lane	*Distribution Type*	*Mean Lane Volume*	*Standard Deviation*
1	Normal	14	5
2	Poisson	8	8

Table 15 - Lane Detail

The results of this are displayed in table 16 shown below.

Lane 1		Lane 2	
Number of Loads	Probability	Number of Loads	Probability
1	99.53%	1	99.70%
2	99.18%	2	98.62%
3	98.61%	3	95.76%
4	97.72%	4	90.04%
5	96.41%	5	80.88%
6	94.52%	6	68.66%
7	91.92%	7	54.70%
8	88.49%	8	40.75%
9	84.13%	9	28.34%
10	78.81%	10	18.41%
11	72.57%	11	11.19%
12	65.54%	12	6.38%
13	57.93%	13	3.42%
14	50.00%	14	1.73%
15	42.07%	15	0.82%
16	34.46%	16	0.37%
17	27.43%	17	0.16%

Table 16 - Tour Reliability

This thesis makes an assumption that the lane volumes are independent of each other, meaning that there is no relationship between the lanes in terms of volume. A high volume on lane 1 does not correlate to a high or low volume on lane 2 in the same period.

Using this assumption, we can determine the reliability of a tour that is made up of these two lanes actually forming for each load volume in question. This is done by simply multiplying the cumulative probability for each value. The result is the probability that all of the tours up to the specified value will, in fact, appear. This is shown in Figure 31. An 'S' shape curve is produced that indicates a relatively low level of risk at high confidence levels followed by a period of substantial risk, denoted by a relatively steep curve. This is followed by a flattening out of the curve as the probability of a tour approaches 0.

Tour Volume	Tour Probability
1	99.23%
2	97.82%
3	94.43%
4	87.99%
5	77.97%
6	64.90%
7	50.29%
8	36.06%
9	23.84%
10	14.51%
11	8.12%
12	4.18%
13	1.98%
14	0.86%
15	0.35%
16	0.13%
17	0.04%

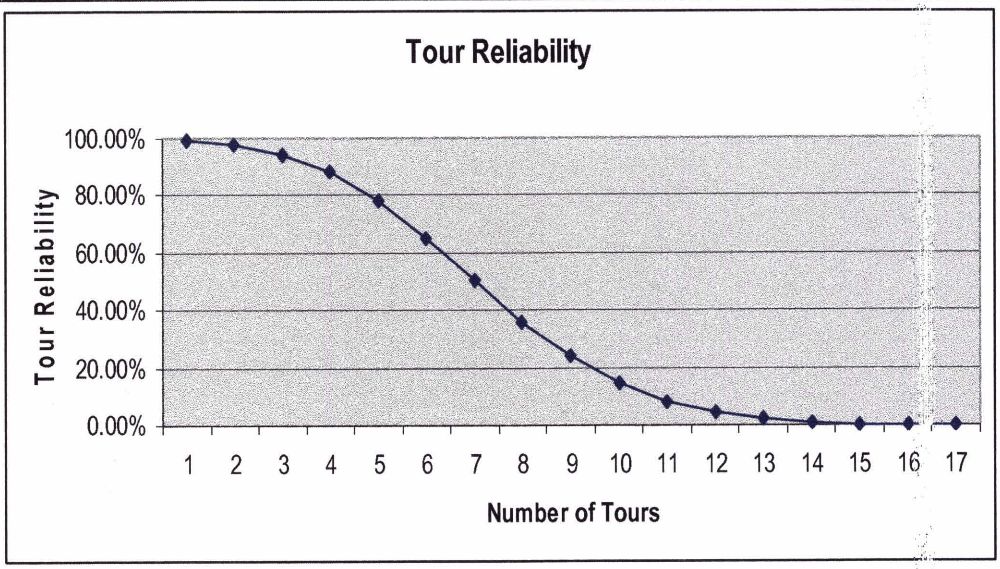

Figure 29 - Tour Reliability and Confidence Levels

Initially, when using a conservative approach represented by a high confidence level, there is a very high probability that the required lane volumes and associated tours will form. This produces very high vehicle utilization, but at the cost of lost potential savings as more expensive for-hire carriers are used.

As more risk is taken by lowering the confidence level, there is an ever increasing probability that fleet capacity will become under utilized. At some point, lower fleet utilization will drive costs up to the point where For-Hire carriers are more attractive for all subsequent volume on that lane.

4.5.5 – Risk Pooling Fleet Assets

Up until now, the simulation process has assumed that a set number of vehicles are assigned by the optimizer to each distinct tour. If one or more loads required to create the tour does not materialize, the fleet asset is left idle and For-Hire carriers are used to execute the load.

For example, assume Tour A has 8 vehicles assigned to it in the optimization, but only 7 tours appear in a given week. Tour B has 2 vehicles assigned to it, but has enough volume to accommodate 3 vehicles. Up until this point, the tour assignment and simulation logic does not allow Tour A to give up its excess capacity to Tour B for that week, however, it is likely that the fleet manager would allocate the unused resource normally assigned to Tour A to Tour B instead of having the asset sit idle.

A more likely use of an internally managed fleet would be to size the fleet based on the forecasted volumes and confidence levels described above. However during operational execution, the fleet would be assigned loads not based on a pre-allocated number of assets dedicated to each specific tour, but based on the optimal use of resources based on each week's realized volume. While the fleet size would remain fixed at each origin facility, each week, the fleet would be used differently based on the actual loads that occurred.

The primary advantage of this technique is that by sharing fleet resources across tours, we enable the ability to take advantage of the likelihood that excess capacity on one tour will be offset by a need for capacity on another tour. Additionally, by aggregating volumes up to the facility level, we are able to increase forecast accuracy and reduce the overall variability of the solution. As noted in chapter 3, when analyzing the CPG company's lanes, even high volume lanes, which had the lowest coefficient of variance, were significantly more volatile than the aggregated volumes of all loads across the network. This increases the overall utilization of the fleet, which ultimately reduces the total solution cost.

To provide an illustration of how risk pooling can be applied, we can look at the current network.

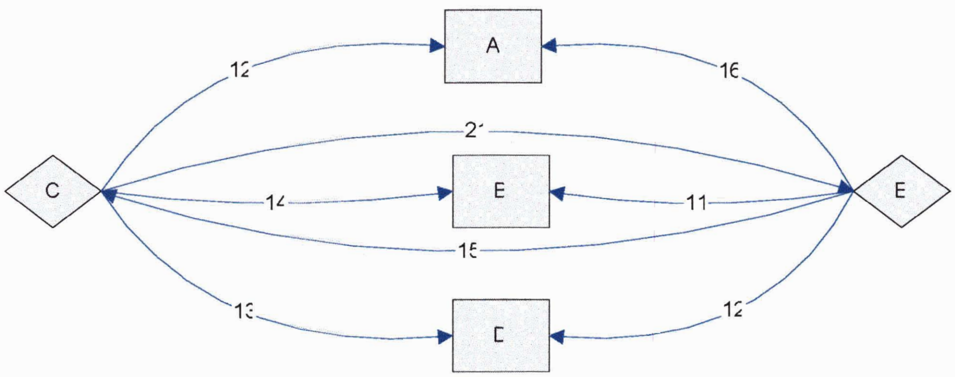

Figure 30 - Network Overview

In this network, we must forecast the volume for each lane segment. Let us assume that the coefficient of variation (C.V.) for each of the lanes is equal to .25. If, however, we aggregate demand back to the origin point, we are able to reduce the overall variability as defined by the C.V. to .125 for all loads originating out of origin point C.

This can be shown by recalculating the standard deviation and mean in the risk pool model.

Mean Loads/Week	Standard Deviation
12	3
21	5.25
14	3.5
13	3.25
60	15

Table 17- Disaggregated Mean and Standard Deviation

The new C.V. is arrived at in the following manner:

$$\text{C.V.} = \frac{(\frac{\sigma}{\sqrt{n}})}{\mu}$$

This transforms to $.125 = \dfrac{\frac{15}{\sqrt{4}}}{60}$

In order to take advantage of risk pooling, the process is to initially optimize the solution as described above. After the initial optimization, instead of simulating demand and constraining the solution based on the fleet capacity allocated to each distinct tour, as has previously been done, a subsequent optimization run is executed using simulated weekly demand. The second optimization run is identical to the initial optimization however it is constrained by the number of fleet assets assigned in the initial optimization. This optimization run does not have the ability to use more assets than were assigned in the initial solve, and solutions that do not utilize all available fleet capacity are penalized for all fleet capacity, as defined in tour duration days, that is left unused.

This allows the optimizer to assign the best tours for that given data set constrained only by the number of available vehicles the optimizer has at is disposal. Specifying the available fleet capacity in the second run is necessary in order to simulate the long-term nature of physical assets, whereby flexing capacity is not a week by week decision, but something made at a higher, strategic level. The results of risk pooling are shown in figure 33. Here we are able to compare the results when each tour is constrained versus when the solution is flexible enough to share resources across tours.

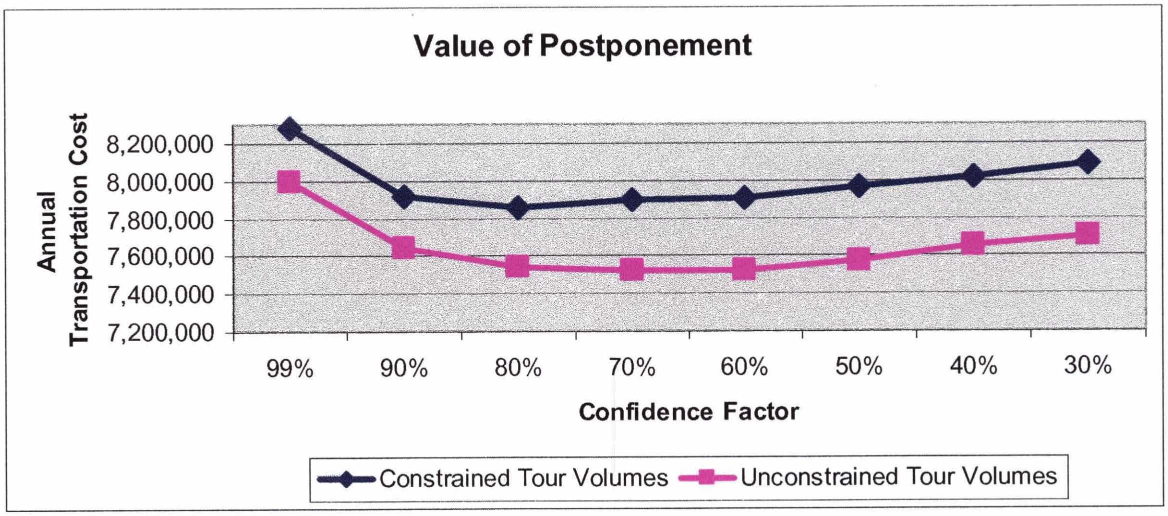

Figure 31 - The Affects of Risk Pooling

The same data can be viewed numerically in table 18. Here we see the savings that risk pooling achieves in relation to the savings generated when fleet resources are tied to a specific tour.

Confidence	Savings (Unconstrained vs. Constrained)
99%	3.42%
90%	3.39%
80%	3.88%
70%	4.77%
60%	4.83%
50%	4.90%
40%	4.44%
30%	4.81%

Table 18 – Savings of Risk Pooling versus Constrained Tour Volumes

When risk pooling is applied, the overall solution costs are reduced at all confidence levels. This is due to two factors:

1. Fleet utilization increases as lanes that are under the forecasted volume are off-set by lanes with higher volume. The re-optimization of the solution enables the fleet to take advantage of these imbalances

2. Additional fleet resources are able to maximize savings by allocating additional fleet resources on the tours that generate the most savings.

It is also important to note that the minimum cost solution was achieved when the confidence level is at approximately 60%, which is markedly below the 80% value that occurs when fleet capacity is tied to specific tours. Still, optimality is arrived at through the use of confidence levels that differ from the mean.

4.5.6 – Utilizing the Private Fleet for Third Party Backhauls

A significant portion of U.S. private fleets are already operating with common carrier authority. This gives shippers the ability to place their internal fleet capacity available on the open market. While some organizations resist the use of their fleet for anything but internal purposes, some shippers do haul third party freight under certain circumstances. This research will not discuss the merits or pitfalls of operating the private fleet as a For-Hire carrier, but will present findings that show how revenue generation on empty lanes impacts the tour formation process.

In this sample, the confidence level used was 90% with a lane C.V. of .25. Figure 34 shows the results of the optimization as the probability of a third party backhaul was increased. In this example, the total transportation cost remained relatively constant at approximately $8M. However the revenue generated by the fleet caused the total transportation costs to decrease when subtracting the anticipated third party freight revenue from the overall transportation cost.

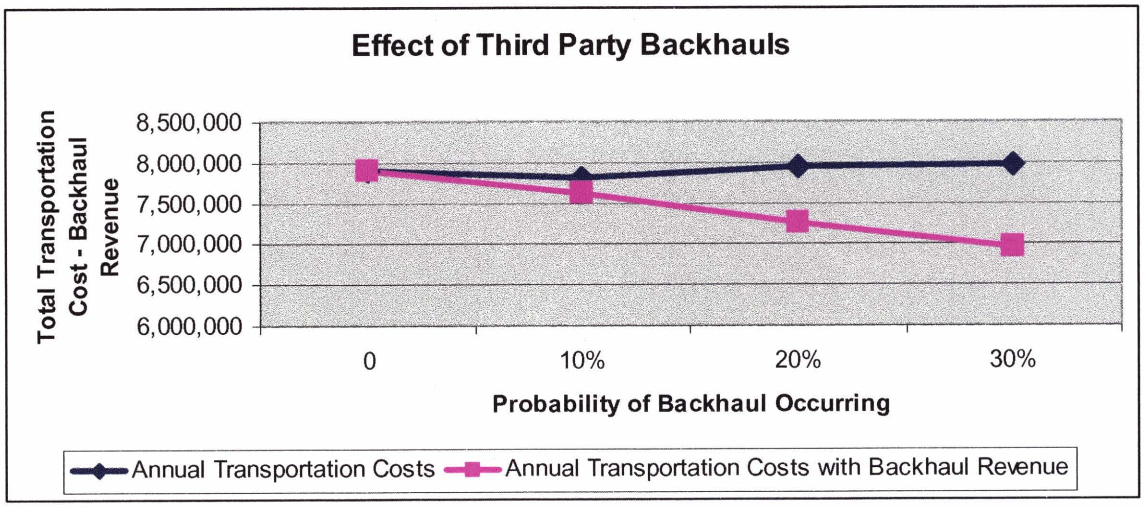

Figure 32 - The Affect of Third Party Freight

The methodology used to capture the third party backhaul revenue was to take the product of the probability of a backhaul occurring times the empty lanes market cost (Rate * Distance). For example, assume there are 10 loads on a lane, each rated at $1000 per load on the open market with a 10% probability of a backhaul. The heuristic will capture $1000 (10% of 10 loads at $1000/load) of backhaul revenue.

An assumption built into the heuristic assumes that the probability of getting a third party backhaul is not related to the number of loads on the lane, but is more of an operational issue that can be represented by a single probability that is used regardless of the number of fleet loads that are allocated to that lane.

Additionally, the heuristic only utilizes third party loads on empty lanes. A case could be made that moving third party freight is a way to increase the number of vehicles allocated to a tour, however this was not incorporated into the heuristic.

The key insight derived from the use of backhauls as part of an overall transportation policy was the relatively low percentage of backhauls required to make tour formation profitable for the fleet. Without backhauls, the scenario in question resulted in 16 trucks and 36 tours. By utilizing a 20% probability of a backhaul on the empty lanes, all available tours became profitable, yielding 25 trucks and 52 tours. These results show the potential of incorporating third party freight into a comprehensive transportation policy.

5 – Conclusions

Long-term trends in the U.S. freight will market reward shippers that have efficient transportation policies in place and penalize shippers who do not. Creating transportation policy that relies on both a private/dedicated fleet we well as contract carriers must be developed in a manner that:

a. Maximizes the shipper's internal economies of scope as well as those of the carrier where possible.
b. Is developed in a manner that is cognizant of the network's normal variability. The use of deterministic methods reliant on weekly or monthly averages to develop transportation policy, as is prevalent today, will cause significant sub-optimality that tends to increase as the variability within a network increases.
c. Utilizes risk pooling as a means to reduce variability and increase utilization of fleet resources
d. Understands the risk/rewards associated with using the fleet to haul third party freight.

Our findings show that for planning, one should use a variability adjusted volume for lanes. While the optimal confidence level will vary for each shipper based on the ratio of the fleet cost to that of for-hire carriers, confidence levels of between .6 and .8 consistently yielded the optimal solution.

Additionally, risk pooling has the potential to significantly reduce costs versus the strict assignment of fleet assets to specific tours. The finds show the potential of a 3-5% savings when risk pooling is employed. However, the optimization did make an assumption that all weekly lane volume was known simultaneously which enabled batch optimization, therefore, the percentage savings shown here does represent an upper bound.

Finally, the use of internal assets to haul third party freight does have economic value. While it is conceded that it may be difficult to execute in real-world operations, the findings show that a relatively moderate use of third party backhauls increases the attractiveness of tours greatly during the tour assignment process. In the study, 100% of possible lane volume was allocated to fleets when just 20% of the empty legs were assumed to haul third party freight.

5.1 –Opportunities for Additional Research

Outlined below are areas that would be logical extensions this thesis.

5.1.1 – Operational Execution

Adhering to the strategic and tactical decisions that were made during the in the development of the transportation policy is a fundamentally different approach then dynamically assigning equipment or carriers based on a snap-shot of information. The dynamic, batch approach is the basis for how most legacy-TMS systems function. These systems often make "greedy" decisions in terms of load assignments because they are unable to predict what loads are likely to arrive.

A methodology that is able to use a probalistic for load assignment would be a welcome addition to the existing literature.

5.1.2 – Quantifying the qualitative side of Carrier vs. Fleet

The reasoning given for using private fleets instead of For-Hire carriers are that they provide shippers with more control, better service and guaranteed capacity. However, can a cost be put on these subjective measures in a way to enable optimization to weigh the trade off. Research better understanding the qualitative reasoning for the use of fleets would enable heuristics to create a more operationally feasible transportation policy using factors above and beyond simply pricing.

5.1.3 – Transportation Forecasting

Determining the anticipated lane volumes utilizing a robust forecasting methodology would be a valuable addition to the literature. The ability to understand the correlation of demand between lanes would enable the creation of more accurate transportation policy, as would be a method to incorporate Sales and Operations Planning outputs into the forecast.

6 – Bibliography

American Trucking Association.2005. The U.S. Truck Driver Shortage: Analysis and Forecast. May, 2005

Arunapuram, S., 1993. Vehicle Routing and Scheduling with Full Loads, Ph.D. Thesis, Department of Operations Research, Case Western University

Bradley, P., 2005. Dude, Where's My Truck, *DC Velocity, v 3, no 5 (May, 2005) pp 48-53*

Bradley, P., 2005. Taking the Company Private, *DC Velocity, v 3, no 8 (August, 2005) pp 44-46*

Caplice, C., 1996. An Optimization Based Bidding Process: A New Framework for Shipper-Carrier Relationships, Ph.D. Thesis, School of Engineering, MIT.

Hall, Randolph W, Racer, Michael., 1995. Transportation with common carrier and private fleets: System assignment and shipment frequency optimization, *IIE Transactions, v 27, no 2 (April) pp 217-225.*

Harding, M., 2005. Can Shippers and Carriers Benefit from More Robust Transportation Planning Methodologies, Masters of Engineering in Logistics Thesis, School of Engineering, MIT.

Johnson, James C, Schneider, Kenneth C., 1988. Private Trucking a Dinosaur, *Business Horizons, v 31, no 1 (Jan/Feb) pp 73-78.*

Livingstone, M, 2006. Senior Vice President of Transportation, C&S Wholesale Grocers,

Lomax, T., Shrank, D., 2005. Urban Mobility Report, Texas Transportation Institute, May 2005

Mulqueen, M., 2006. Personal experience as well as conversations with various industry experts from shipper, carrier, and third party logistics firms that elect to remain anonymous

Powell, Warren B., Sheffi, Yosef, Nickerson, Kenneth S., Butterbaugh, Kevin, Atherton, Susan., 1988. Maximizing Profits for North American Van Lines' Truckload Division: A New Framework for Pricing and Operations, *Interfaces, v 18, no 1 (Jan/Feb) pp 21-42.*

Ronen, D., 1997. Alternate Mode Dispatching: The Impact of Cost Minimization, *The Journal of the Operational Research Society, v 48, no 10 (Oct) pp 973-977.*

Standard & Poor's Corporation, 2006. Industry Surveys, Transportation: Commercial. (February 9, 2006)

Terreri, A., 2006. Shippers are Adding Private Fleets to their Transportation Mix, *World Trade Magazine,* on-line edition Feb 1, 2006

Transportation Topics, 2005. Transportation Topics Top 100 Private Fleets, (2005)