Characteristics of Spot-market Rate Indexes for Truckload Transportation

by

Andrew Bignell

B.A. (Hons), Business Administration
University of Western Ontario, 2007

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 2013

Signature of Author..

Andrew Bignell
Master of Engineering in Logistics Program, Engineering Systems Division
May 30, 2013

Certified by..

Dr. Chris Caplice
Executive Director, Center for Transportation and Logistics
Thesis Supervisor

Accepted by.. ....

Prof Yossi Sheffi
Elisha Gray II Professor, Engineering Systems Division
Professor, Civil and Environmental Engineering Department
Director, Center for Transportation and Logistics
Director, Engineering Systems Division

Characteristics of Spot-market Rate Indexes for Truckload Transportation

by

Andrew Bignell

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

ABSTRACT

In the truckload transportation industry in the United States, a number of indexes are published that attempt to measure changes in rates, but no single index has emerged as an industry standard. Industry participants, particularly those exposed to the spot-market, have found that existing indexes do not effectively represent their experiences in the marketplace. A widely-used and valid spot-market index could allow for more effective tactical decision-making, the development of freight derivatives, stronger analysis and negotiation of contract rates, and contracts with index-tied rates. This paper examines pricing indexes from other industries and expands upon a framework of characteristics that support successful indexes. Using this framework to evaluate existing industry indexes, it finds that two commonly used indexes are not designed appropriately for use by the spot-market. It also examines rates experienced by a large North American provider of non-asset based logistics services and finds that they differ significantly from rates measured by existing indexes. The analysis suggests that indexes of the spot market would be improved by disaggregating rate information based on geography and tender lead time.

Thesis Supervisor: Dr. Chris Caplice
Title: Executive Director, Center for Transportation and Logistics

Acknowledgements

To my MIT Supply Chain Management colleagues who have always been free with their time and ideas as this thesis has taken shape.

And also to Dr. Chris Caplice and Thea Singer, for their ongoing support and advice.

Finally, to my parents - they always believe in me and support me in my ongoing adventures.

Thank you.

Table of contents

List of figures

List of tables

1. Introduction

Indexes are used across many areas of the economy. They are used to measure the stock market, economic growth, price changes and even social measures like journalistic freedom. While many are tracked only by specialists in their field, some, like the Consumer Price Index (CPI), are well-known to the general public and influence economic policy decision-making at the highest level. This thesis examines the important questions and considerations of index design for the truckload transportation spot-market in the United States.

An index, simply, is a statistical measure of changes over time in a representative set of data points. Commonly, an index is described using index numbers. A base period and value (often 100) is defined and the value of the index at any given time is compared to that base level. For example, a price index level of 150 would indicate that prices are now 50% higher than prices observed in the base period.

Indexes can be used for a number of purposes. They provide a simplified view of complex data that can help decision-making. Often index values are used as inputs in other calculations. The S&P 500 stock index, a measure of the price level of 500 leading US companies, is a component in another index, the Composite Index of Leading Indicators, that provides insight into the overall direction of the economy. Indexes are also used as inputs in contracted agreements. CPI and other indexes of general price levels are used to calculate the amount owed to retirees with defined benefit pension plans.

Indexes are calculated in different ways according to their own stated purpose. Some indexes are intended to represent only the data within the sample. The S&P 500 Index, for example, only measures changes in the price of its 500 component stocks. Even though some analysts do believe its movements correlate with broader economic conditions, the index is not designed, or intended, to represent anything beyond its sample data. Alternatively, the intent of the CPI is to measure changes in prices across all consumer goods in the economy. It would be impossible to measure all prices, so the CPI determines a representative sample of goods and collects data only from that sample. The challenge with an index of this type is to ensure that the sample does reflect the market as a whole and that the sampling method does not insert any biases.

This chapter reviews the characteristics of the truckload industry and discusses how changes in spot-market rates can impact operations of shippers, carriers and brokers. It also examines the factors that drive changes in rates.

1.1 Trucking industry basics

Trucking is the predominant mode of commercial freight transportation in the United States. In 2010, the industry had revenues of $563 billion and claimed 81% of the overall commercial transport market (Standard & Poors, 2012). Despite rising fuel costs and increased capacity of intermodal railway providers, this dominance has only declined marginally from 2004 when trucking captured nearly 87% of the market (Standard & Poors, 2004).

The trucking industry can be segmented into two broad categories - private carriage and for-hire carriage. Private carriage – when a company uses their own fleet of tractors to transport their goods – accounts for an estimated $268 billion (48% of the trucking industry). For-hire carriage – when a company contracts with a transportation provider - makes up the remainder.

The for-hire trucking industry can be further segmented into full-truckload (TL) and less-than-truckload (LTL). A TL operation involves freight from only a single shipper and generally is moved directly between a single origin and destination. Less commonly, a TL carrier may make multiple stops to pick-up or drop-off goods if the shipper has different physical locations that require service. An estimated

45,000 carriers serve the TL market and the industry is highly fragmented. Approximately 30,000 carriers have revenues less than $1 million and the largest carrier in 2011, J.B. Hunt, had revenues of $4.5 billion and a market share of less than 2% (Standard & Poors, 2012). By contrast, an LTL carrier picks up many, relatively small loads from multiple customers and consolidates them at one of their hub facilities. The LTL industry is more concentrated and the largest provider, Con-way Transport, had revenues of $5.3 billion and a market share of 13% in 2011 (Standard & Poors, 2012). The research in this paper is confined to TL operations.

There are two main parties in trucking transportation: the shipper, and the carrier. The shipper has goods that require transportation. The carrier has assets – trucks and trailers – that can be used in the process of transporting those goods from their origin to their destination. Shippers may contract directly with carriers or they may contract through an intermediary. The two main types of intermediaries are third-party logistics providers and freight brokers.

Third-party logistics providers (3PLs) serve shippers that wish to outsource some or all of the processes involved in transportation management. In this role, 3PLs will negotiate rates with carriers and manage day-to-day operations for their clients. Many carriers have also developed dedicated departments to provide 3PL services to shippers.

Freight brokers initially arose to sell, for a commission, the unused capacity of carriers that had a sales force with limited capabilities or geographic reach. Over time, shippers also began to hire brokers to find carriers when they needed transportation. Today, freight brokers often enter into long-term agreements with shippers to provide transportation on some or all of a shipper's lanes. The broker then hires carriers as needed to supply transportation for the shipper. In this capacity, brokers are often referred to as non-asset-based carriers.

1.2 Procurement of transportation

Shippers procure transportation services from for-hire carriers in one of two ways. The most common, contracting, establishes a guaranteed rate for service provided by the carrier for a period of time – usually 1-2 years. Rates are established either through negotiation or an auction process. Shippers may also access the spot market. A spot transaction is usually for a single load and the price is determined at the time of booking. Although exact figures are difficult to determine, the spot market is estimated to comprise less than 10% of the total TL market (Caplice, 2007). This value is supported by an Armstrong & Associates listing of the revenues of the thirty largest freight brokerages. In 2010, their combined revenue was $22.2 billion, accounting for approximately 8% of the total truckload market.

1.2.1 The spot market

The spot market is a catch-all term for any transaction where the price is determined at (or near) the time when the load is actually moved. Shippers may offer a load for tender to their carrier base through their Transportation Management System (TMS). However, the shipper is limited by the number of carriers maintained in the TMS and may not receive a satisfactory rate. To access a larger pool of carriers, shippers may contract with a freight broker.

Freight brokers maintain relationships with thousands of carriers and shippers and are able to match the supply and demand for transportation. As of 2010, there are approximately 16,000 registered freight brokers in the United States (Supply Chain Digest, 2010). They range in size from offices of one or two people to sophisticated enterprises with multiple facilities spread across the country and thousands of employees. The largest, C.H. Robinson Worldwide, had revenues of $9.3 billion in 2010 and accounted for slightly more than 40% of the revenues of the top thirty largest brokerages (Armstrong & Associates, 2012). Brokers have trained staff and operational systems that allow them to quickly determine which carriers have capacity available and they negotiate strongly to achieve the lowest price.

A brokered spot-market transaction begins with a negotiation between the shipper and a freight broker. A shipper may request prices from several brokers and the market is often highly competitive. Once a shipper finalizes a price with a broker, the broker will begin work to find an available carrier within their network of contacts. Carriers may also provide information to brokers about where they have vehicles and drivers available. Sometimes a match is easily made; at other times, booking a carrier may require dozens of phone calls. Brokers make a margin between the rate they have established with the shipper and the cost they are able to negotiate with the carrier. Because rates to shippers are established prior to negotiations with carriers, the margin achieved by a broker on any particular transaction can vary widely and may even be negative.

1.2.2 Contracting with non-asset based carriers

An additional procurement system has developed that bridges the gap between contracting and the spot market. Non-asset based carriers (freight brokers) participate in auctions and provide rates to shippers in the hopes of being awarded steady business on some lanes. If they are awarded lanes, the brokers will then depend on their network of asset-based carriers to physically provide the service. Occasionally, the freight brokers negotiate contracted long-term rates with one or more carriers. More frequently, they will use their spot market operations to negotiate with carriers the cost of each individual load while receiving a set rate from the shipper.

The main advantage of freight brokers is their ability to flex capacity easily within their network of carriers. Successful brokers also manage many of the other service concerns of shippers. They prescreen carriers for reliability and, through online tools, can provide updates and information to both carrier and shipper in ways that mirror the services provided by an asset-based carrier.

Freight brokers that contract business in this way are making a bet that the aggregate amount that they pay to their carriers over the course of the contract will be less than the aggregate revenue received from the shipper. Their profitability depends heavily on the prevailing spot market rates.

1.3 Rate variability

Throughout this paper, rate variability is defined as variability within the spot market only. Because most industry players (both shippers and carriers) are exposed to, at least some, spot market transactions, rate variability can have a direct impact on their costs or revenues. However, rate variability will also tend to have broader short-term and long-term impacts. Section 1.3.1 discusses how high spot-market rates increase the probability of carriers rejecting tenders from shippers and section 1.3.2 discusses how rate variability impacts operations of brokerages.

1.3.1 Impact on contract execution

As noted previously, contracts in truckload transportation guarantee a price, but do not necessarily guarantee the availability of the equipment. While the rate of contract failure will vary based on the characteristics of the carrier, it is also likely to increase when rates are high in the spot-market. This results from both indirect and direct mechanisms.

The indirect mechanism occurs because both spot-market rates and rejection of tenders by carriers are indications of market tightness (when demand temporarily exceeds supply). Most carriers cannot afford to reserve capacity for any particular customer. Therefore, as demand for transportation increases, the carrier will reach its capacity limits sooner. Once a carrier fills its entire capacity, it will start turning down tenders from some customers.

Increased rate of contract failure can also result directly from the opportunistic behavior of carriers. Most carriers will look to the spot market to fill capacity during slower periods. They may also look to the spot market to maximize revenues during periods of high demand. During periods of market tightness, it is very likely that the rate a carrier could achieve in the spot market will exceed the rates they have agreed with their contract customers. Therefore, the carriers will face an incentive to increase their exposure to the spot market during periods of high demand and reduce the capacity they have available to serve their contracted customers. Shippers may face increased tender rejection as carriers intentionally move capacity to higher value opportunities.

When a shipper's primary carrier rejects a tender, the shipper will usually face increased costs. Research has shown that each initial tender rejection increased costs by 13.9% (Kim, 2013).

1.3.2 Impact on brokerages

Brokerages are usually fully exposed to spot-market pricing. In a simple transaction for a single load, the broker will provide a price quote to the shipper. The broker makes their pricing decision prior to securing capacity through an asset-based carrier, and therefore, the price is based on an expectation of their costs. Expectations are based on recent experiences, knowledge of events or conditions in the marketplace and the particular characteristics of the load. Actual costs incurred can vary widely from expectations and, consequently, achieved margins can be negatively impacted.

Brokerages can be impacted even more significantly when they act as a non-asset-based carrier. In this role, they establish a contracted rate for a period of time with a shipper. Again, the contracted rate offered by the brokerage will be based on the expectations of future spot-market prices and their desired margin. However, as this contract lasts for a period of time and may cover thousands of transactions, the impact of incorrect expectations on profitability is much greater.

1.4 Factors that influence spot-market rates

It is important to understand the drivers of transportation rates when discussing transportation indexes. The rates experienced in the market are impacted by a variety of factors and any index must choose how to account for these factors. Two possible approaches are normalizing data to remove the effects of some factors and specifying the index in such a way that some data is automatically excluded.

The vast majority of research evaluating drivers of transportation rates has focused on contracted rates. This research has investigated actions that a particular shipper may take that could have positive or negative impacts on their total cost of procuring transportation. However, the spot market for transportation may not be impacted in the same way by some variables.

1.4.1 Geographic factors

A large portion of the costs of providing transportation are driven by geographic factors. The distance that a load travels determines the amount of fuel used, the wear-and-tear on the tires and trucks, and, since drivers are paid by the mile, the labor cost incurred.

However, distance is not the sole geographic factor. Each combination of specific origin and destination is called a lane. Carriers will tend to view certain lanes more favorably than others and will offer lower rates on more attractive lanes. Generally, carriers will provide lower rates for high volume lanes where there is a high probability that they will be able to schedule a follow-on pick-up nearby. The ultimate situation, a "continuous move", involves a carrier delivering an inbound load and picking up an outbound load at the same facility on the same day. Carriers will typically discount rates 5-8% for this opportunity as it eliminates deadhead miles and reduces dwell time (Caplice & Sheffi, 2006).

The attractiveness of a lane will depend on both the probability of achieving a nearby pickup and the expected outbound rate. For a contract carrier, the expected probability of picking up a nearby load and the expected rates for that load are dependent on the portfolio of businesses they serve in that market and can be relatively well understood. By contrast, for a spot-market carrier the expected probability of picking up a load and the expected rate are more dynamic and will depend on the current supply and demand in that market. If the supply-demand balance is volatile, these factors can change rapidly.

Carriers may also use current information in their tactical decision-making instead of long-term expectations. As an example, a carrier may choose to commit to pick up a load at Location A three days from now. In choosing loads today, the carrier would look more favorably on a load destined for Location A as it would allow the carrier to reposition their equipment to serve that future business.

Therefore, the attractiveness of a particular spot market lane for a carrier depends on a combination of their short-term business decisions and the dynamic situation in the overall spot market. These geographic factors on rates are likely to be viewed very differently by different carriers. The same carrier may even evaluate the attractiveness of a lane differently on different days. These factors lead to a spot market with significant volatility.

1.4.2 Other factors

There are several other factors that influence rates in the spot-market. Clear pricing differences exist for different equipment types. The two most common equipment types in the trucking industry are standard dry vans (either 48' or 53') and refrigerated vans. Refrigerated vans have attached diesel-fueled cooling units and have higher operating costs than standard vans.

Tender lead time has been shown to have a significant impact on rates (Caldwell & Fisher, 2008). For shippers with contracted rates, tender rejection increased steadily as lead time declined below five days. Average rates rose as the shipper was forced to go deeper into their routing guide. Caldwell & Fisher found that shippers providing three days of lead time had rates that were 4% lower than shippers that provided only one day of lead time. This relationship appears to hold even more strongly in the spot-market and is discussed further in Chapter 6.

Caldwell & Fisher (2008) also identified Day of the Week as a minor driver of rate differences. They noted that costs for Wednesday-Friday were higher than Monday-Tuesday and speculated that it was due to standard five day work weeks that tended to constrain capacity later in the week.

1.4.3 Fuel

Fuel is also a direct driver of cost. Unlike other factors, fuel costs are usually explicitly accounted for in contracts through the use of fuel surcharge programs. Fuel surcharges are agreed premiums to be paid by the shipper to the carrier based on the cost of fuel at the time the load is transported. Fuel surcharge programs use a diesel fuel price index published by the Department of Energy to estimate the current fuel price. Some surcharge programs use the national diesel fuel index, while others will use one of several regional diesel fuel indexes. In the spot market, however, fuel surcharge programs may or may not be used. Sometimes rates are negotiated as an "all-in" price. An index that measures the spot-market will have to determine whether it should be inclusive or exclusive of fuel and manage its data appropriately.

1.5 Research questions

This paper discusses and analyzes the potential development and use of indexes of the truckload spot-market. Comparisons are also made to indexes in other industries. It develops a checklist of index characteristics that would best support widespread adoption by the industry. This paper surveys existing indexes and determine which, if any, possess the characteristics identified in the checklist.

This paper will also provide a more tactical discussion of the key questions that need to be answered in the design of an index for the truckload spot-market. These question include: 1) what units should the index use, 2) how often should the index be updated, 3) where should data be collected, 4) what data, if any, should be excluded, 5) should data be normalized across a variety of factors and, if so, how, 6) should data be adjusted for seasonal factors and, if so, how, 7) should sub-indexes be used to provide more granular information.

Finally, this paper will address some of the structural questions about who should manage the index and how data should be accessed and shared. These non-technical considerations are equally important in the design of an index that will actually be used and adopted by the industry.

1.6 Chapter summary

This chapter introduced the truckload transportation industry in the United States with particular focus on how variability of rates in the market can cause operational challenges for shippers, carriers and brokerages. It also discussed the key factors that influence rates for the spot-market.

1.7 Looking ahead

Chapter 2 examines successful indexes from other industries and develops a framework that can be used to evaluate the design and management structure of indexes. Chapter 3 uses this framework in the context of the rate indexes that currently exist in the US trucking industry. Chapter 4 examines a dataset of spot market transactions provided by a major freight brokerage and compares observed rates with existing industry rate indexes. Chapter 5 discusses the differences in rate behavior that can be observed in different geographies. Chapter 6 investigates the relationship between lead time and rates and discusses how this relationship changes according to market conditions. Chapter 7 uses industry data to construct a lane index. Chapter 8 measures the accuracy of indexes and discusses the viability of rate indexes for a number of potential uses. Chapter 9 reviews the implications of this research and suggests topics for further research.

2. Indexes

This chapter discusses two important aspects of indexes for the spot truckload transportation market. Section 2.1 examines the broad considerations of what an index should look like. To do this, a set of characteristics developed from the maritime industry are examined and expanded upon. Section 2.2 discusses the potential uses of an index of the spot truckload market.

2.1 Index design characteristics – basic framework

Because the uses of indexes are often particular to the needs of each individual market, there is no definitive set of rules that govern how indexes should be designed and maintained. A number of indexes and pricing benchmarks have been developed for commodities, like forest products and fisheries, where prices are not transparent. Much of the trade in these products occurs outside major exchanges. These commodities share some of the characteristics of truckload transportation and the development and use of their price indexes provide some insight.

The industry most closely resembling the truckload transportation industry is the maritime shipping industry. The primary indexes in maritime freight are published by the Baltic Exchange of London. The Baltic Exchange can trace its history back to 1744, but their experience in managing indexes began in January 1985. The objective of their first index, the Baltic Freight Index (BFI) was two-fold; to be a "barometer" of worldwide freight markets and to provide an underlying commodity for freight derivatives (Kavussanos & Visvikis, 2006). Because spot transportation markets in the trucking industry and in the maritime industry share similar fundamental structures, the design and history of the Baltic Exchange indexes can provide lessons for the development of trucking indexes.

Kavussanos & Visvikis (2006), use a previous discussion by SSY Futures, an international shipbroking and analysis firm, to describe a set of ten characteristics that are necessary for freight indexes. However, after examining the structure and history of existing indexes, it was found that the framework could be improved with the inclusion of two additional characteristics. Characteristics #6: Disaggregation and #12: Updating are new additions to the framework. The evidence for the inclusion of these characteristics is discussed in sections 2.1.1 and 2.1.2. The twelve total characteristics can be logically segmented into two categories. The first category contains those characteristics dealing with the actual calculation of the index. The second category contains those characteristics associated with the management structure of the index. All twelve characteristics are listed in Table 2.1

Table 2.1: Characteristics of freight indexes

Index calculation and data collection:

1. **Accuracy:** Should be an accurate reflection of the real spot market
2. **Bias:** Should be rigorously computed and unbiased
3. **Familiarity:** Should be expressed in units familiar to the industry
4. **Balanced input:** Should be based on a sufficiently broad and balanced input
5. **Transparency:** Should be transparent and simple
6. **Disaggregation:** Should provide different levels of aggregated information in a clear and calculable hierarchy

Index management:

7. **Frequent publication:** Should be published regularly and frequently (preferably daily via electronic distribution)
8. **Auditing:** Should be audited and monitored by an independent body
9. **Complaints:** Should have proper procedures for dealing with complaints
10. **Cost:** Should be low-cost
11. **Participation:** Should be supported by the major participants in the market
12. **Updating:** Should have procedures for updating and adjusting components or index structure as market conditions change

Evidence for the validity of these characteristics can primarily be shown through an examination of the structure and history of the indexes of the Baltic Exchange, but they are also supported by indexes from other industries.

2.1.1 Index calculation and data collection – evidence from other industries

The Baltic Exchange Indexes are calculated by aggregating rates across a number of routes. The initial index was comprised of 13 different routes carried by vessels of varying size with weightings in the index of between 2.5% and 20% **(#5 Transparency)**. The Baltic indexes are calculated as unitless indexes. The first index was established with a baseline of 1000 in 1985 **(#3 Familiarity)**. Table 2.2 shows the initial composition of the index in 1985.

Table 2.2: The Baltic Freight Index, January 1985 (adapted from Kavussanos & Visivikis, 2006)

Routes	Vessel Size (dwt)	Cargo	Route Description	Weights
1	55,000	Light cargo	US Gulf to Netherlands	20%
2	52,000	Heavy Grain	US Gulf to South Japan	20%
3	52,000	Heavy Grain	US Pacific Coast to South Japan	15%
4	21,000	Heavy Grain	Us Gulf to Venezuela	5%
5	20,000	Barley	Antwerp to Red Sea	5%
6	120,000	Coal	Hampton Roads (US) to South Japan	5%
7	65,000	Coal	Hampton Roads (US) to Netherlands	5%
8	110,000	Coal	Queensland to Netherlands	5%
9	55,000	Coal	Vancouver to Netherlands	5%
10	90,000	Iron Ore	Monrovia to Netherlands	5%
11	20,000	Sugar	Recife to US East Coast	5%
12	20,000	Potash	Hamburg to West Coast India	2.5%
13	14,000	Phosphate	Aqaba, Jordan to West Coast India	2.5%

The Baltic Exchange publishes both an overall index and the rates on each of the individual routes that comprise the index. By providing information in this way, market participants can access information that is specific to the routes that are important to them. Over time, pressure from the industry pushed the Baltic Exchange into offering even more specific information. In 1998, they created the Baltic Panamax Index (BPI) that tracked rates for vessels between 55,000 and 75,000 deadweight tons (dwt). In 1999, they added the Baltic Capesize Index (BCI) for vessels between 150,000 and 172,000 dwt. Today, they measure rates on more than 50 routes that form the basis for four major indexes delineated by vessel size.

These four major indexes are then aggregated to calculate the Baltic Dry Index, a measure of overall rates in the maritime freight industry that is published daily and watched worldwide **(#6 Disaggregation)**. This structure is shown below in Figure 2.1.

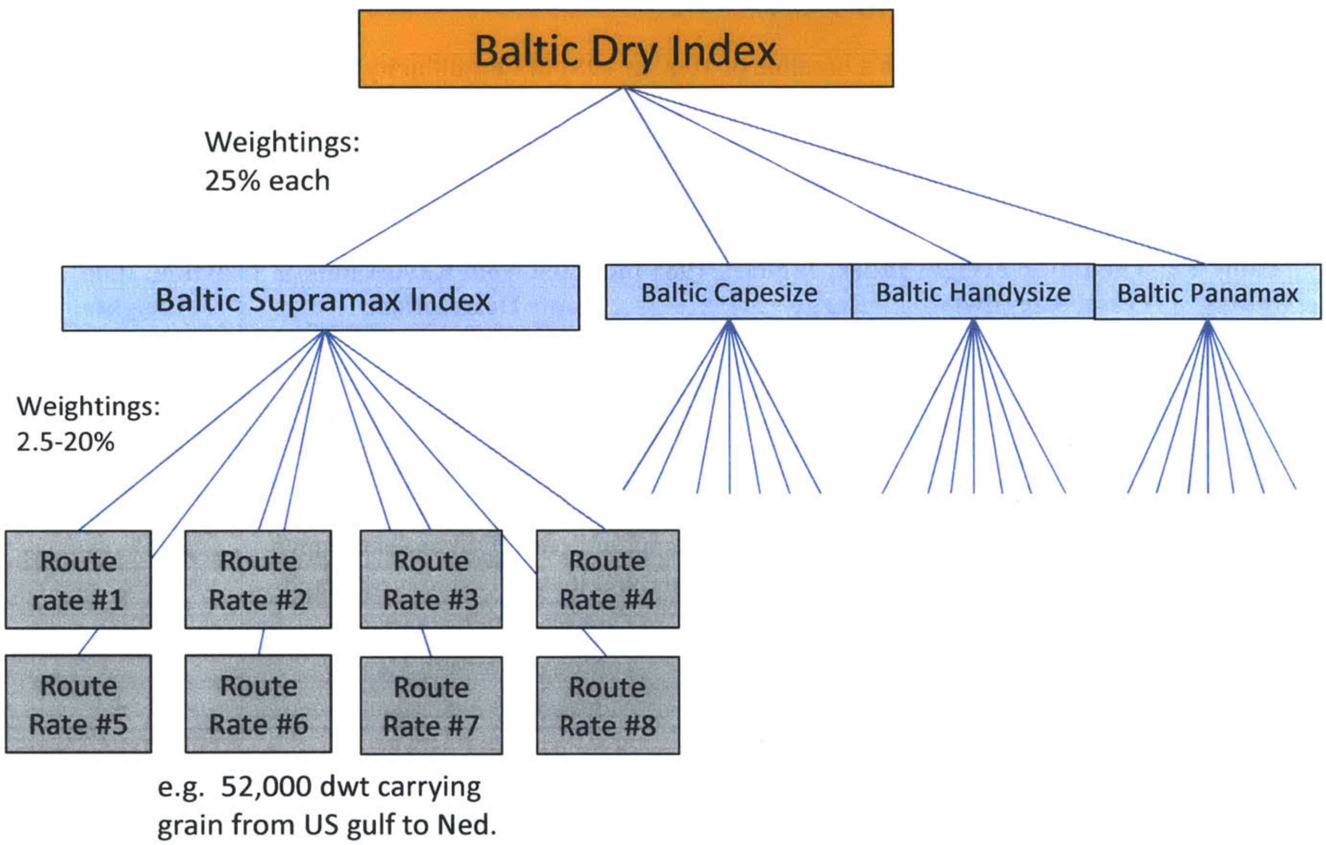

Figure 2.1: Hierarchy structure of the Baltic Dry Index

Determining which routes to include in the indexes is the purview of the Freight Indices and Futures Committee (FIFC), made up of five directors of the exchange. They consult with Baltic Exchange members, user groups, other exchanges and interested parties in their decision process, but are also guided by a standard set of principles governing route selection. The list below, summarized from the Baltic Exchange's *Manual for Panellists,* describes these principles.

- The need to produce a basket of routes which is as representative as possible of the world's principal bulk cargo trades;
- Geographical Balance. Routes reflect trades both within the Pacific and the Atlantic, as well as trades between the oceans (maintaining a balance between fronthaul and backhaul routes);
- Liquidity. A steady and significant turnover of fixtures on index routes, or on routes related to them is important. Trades subject to seasonal closures are avoided;
- Transparency. A reasonable volume of accurately reported fixtures should be available. Where possible, trades dominated by a sole or limited number of charters are avoided;
- Standard Terms. Voyage routes where business is largely concluded on standard terms are favoured;
- Number of routes. The Baltic is conscious of the daily workload of panellists, and of the limited number of trades in the market place which fulfil their selection criteria.

The principles ensure the Baltic Exchange Indexes represent a broad range of routes, all of which have sufficient liquidity to be computed accurately **(#1 Accuracy and #2 Bias)**. The process for calculating the rates on each route is also carefully managed through the participation of a select group of panellists. The panellists provide their best judgement of the prevailing market rates each day. The Baltic Exchange believes that reporting panels are the best option as there is no "independently verifiable 'right' or 'wrong' rate" (Baltic Exchange, 2011.) The Baltic Exchange's *Manual for Panellists* also describes the criteria for selecting panellists and is shown below.

- The main business of the panellists should be shipbroking. Principals are not considered appropriate panellists;
- Panellists must be recognised as competent, professional firms, actively engaged in the markets they report, with adequate personnel to perform the role of panellist;
- Panellists must be members of the Baltic Exchange, fulfilling all relevant membership criteria;
- An appropriate geographical spread of panellists is maintained;
- The Baltic seeks to avoid the appointment of panellists who are the exclusive representatives of charterers who are particularly influential in relevant trades;

The number of panellists for each index ranges from 11 (Baltic Capesize Index) to 17 (Baltic Panamax Index). While panellists represent a small proportion of the overall membership of the Baltic Exchange – approximately 600 members – panellists tend to be from the larger shipping brokerages. The Baltic

Exchange's approach to selecting panellists attempts to ensure a broad and balanced input **(#1 Accuracy and #4 Balanced input)**.

Indexes from the forestry industry also demonstrate the importance of providing disaggregated information **(#6 Disaggregation)**. Pulpex was once one of several index and derivatives providers active in the forestry industry. They focused their efforts on a specific forest product, *Northern Bleached Softwood Kraft (NBSK) delivered to Northern Europe*. Although it is one of the most frequently traded variants of pulp, it may not have correlated well with other variations in different regions around the world. A lack of correlation reduced the usefulness of the index and its derivatives for hedging purposes and Pulpex failed in 2003 (Paul, 2003). By contrast, FOEX, the surviving producer of pricing indexes in the pulp and paper industry has launched approximately 1.5 indexes per year since 1996. Each new index focuses on a specialized portion of the pulp, paper and forestry industry. As an example, in June 2011, they launched the PIX NBSK pulp China index that measures spot pulp pricing in the Chinese market (FOEX, 2013).

Indexes developed for the salmon fishing industry also demonstrate some of these characteristics. In 2007, two separate indexes launched independently. The surviving index, Fishpool, is calculated by measuring prices in five different markets and combining them according to simple, published weightings **(#4 Balanced input and #5 Transparency).** Their failed competitor, FishEx, depended solely on Norwegian export statistics for their index calculation (Gronvik, 2008).

2.1.2 Index management – evidence from other industries

The characteristics of the Baltic Exchange indexes also align with the characteristics of index management from Table 2.1. The Baltic Exchange is published daily **(#7 Frequent publication).** It is supported by a significant fraction of the participants in the maritime shipping industry. The Baltic

Exchange claims that more than half the world's tonnage is traded by Baltic Exchange members **(#11 Participation).** The index is also relatively low-cost. Annual membership fees range from £1040 for a sole trader to £26,215 for a corporate membership for an organization of more than 50 members. Non-members can also access detailed freight information for £3905 per year **(#10 Cost)**

The Baltic Indexes also have clear processes for making changes or adjustments to route assessments and index composition **(#12 Updating)**. The FIFC always consults with the market – their members – when evaluating new products to offer. Once a product has been identified, they evaluate the changes through extensive trial runs. They telegraph any changes months in advance and often they will run two similar indexes concurrently for up to a year to allow traders to exit positions in the original index and transition to the new index.

Fishpool, the surviving salmon fisheries index is published daily **(#7 Frequent publication)** and has no membership fee **(#10 Cost).** Their failed competitor, FishEx, published once per week and had a membership fee of 750 euros.

FOEX, active in the forestry sector provides index data that is audited monthly by Ernst & Young **(#8 Auditing).** They also have established complaint procedures that involve an independent advisory board **(#8 Auditing and #9 Complaints)**.

A cautionary tale of an index that has failed to maintain independent monitoring and auditing of indexes **(#8 Auditing)** is the spectacular LIBOR scandal that engulfed major banks starting in 2008. LIBOR is the *London Interbank Offered Rate* and is intended to measure the rates that banks would have to pay to borrow money from other banks for a certain period of time and in a specific currency. Data is provided by up to 18 financial institutions in London on a daily basis. Thomson Reuters collects the submissions, removes high and low outliers and then publishes the daily LIBOR rate. This rate is then used as a benchmark for roughly $10 trillion in loans and $350 trillion in derivatives. Starting in 2005, Barclays bank employees began to falsify their submissions to Thomson Reuters with the intention of affecting the published daily LIBOR rates. Their objective was to assure a rate that would be beneficial to their firm's trading position. In 2012, Barclays agreed to pay more than $453 million in fines (Gallu, Brush & Fortado, 2012).

2.2 Roles of Indexes

An index of the truckload spot-market could perform several functions. It could provide better visibility of prices to participants in the market. It could also potentially be used as a basis for financial derivatives. A spot-market index could be included as an input into contracted rates between shippers and carriers or between shippers and brokers. It could also be used by shippers to compare the costs of entering into long-term contracts with the costs of acquiring transportation as needed in the spot-market. The remainder of this chapter will discuss each of these uses in more detail.

2.2.1 Price visibility and opportunistic optimization

Price visibility in truckload transportation would open new avenues for optimization by participants. Carriers would be able to increase their exposure to the spot-market and drive incremental revenue during tight-market conditions. Shippers might choose to access the spot-market opportunistically when rates are low and reduce their overall freight spend. In procurement events, shippers with knowledge of the history of spot-market rates would be able to better benchmark expected rates for each lane.

2.2.2 Derivatives

An index that measures spot transportation rates could also potentially allow for the development of various financial derivatives. Simply put, a derivative is a financial asset whose value is dependent on another underlying variable. In this case, the index would act as the underlying variable and several different financial structures could be developed upon this base.

Financial derivatives are useful because they allow for risk management through financial hedging rather than through changes in operations. If the derivatives are designed well and the underlying asset (the index) effectively measures changes in rates, shippers and carriers could use these derivatives to lock-in long term rates without contracts. Managing risk through financial derivatives can often be cheaper as the transaction costs can be low.

Derivatives also open doors for greater involvement in transportation markets by non-operating financial companies. While many traditional operators may be fearful of speculators and market manipulators, the presence of financial participants is common in most commodities. Speculators add liquidity to markets and tend to dampen violent price changes caused by supply or demand shocks. Additionally, evidence from the dry-bulk shipping market shows that prices of freight derivatives incorporate additional information that is not incorporated in spot-market prices (Kavussanos & Visvikis, 2004). This additional information can encourage better medium and long-term decision-making by operators.

2.2.3 Development of risk-sharing contracts

Indexes can also be used to change the nature of contracts. The vast-majority of contracts and transactions in truckload transportation include a fuel surcharge. A contract that involves a fuel surcharge will be comprised of a line-haul charge that is constant throughout the period of the contract and a fuel surcharge that will vary depending on the price of fuel. Both shippers and carriers see advantages in calculating fuel costs separately. The shipper is able to focus on the portion of their transportation costs within their control. The carrier prefers fuel charges to be calculated separately so that they can set long-term rates without needing to account for the volatility of fuel prices. By removing one major source of cost variability, fuel surcharge programs can help ensure carrier solvency. Indexes that track spot-market rates could be used in an analogous way.

A shipper and a carrier may choose to design a contract where a portion of the carrier's revenues is flexible and dependent on an index of spot-market rates. For an asset-based carrier, this would mean increased revenues when the rates are high. This should reduce the incentives for the carrier to reject a shipper's tenders and increase the overall service level through the length of the contract. Similarly, the shipper will benefit from lower costs when spot-market rates are low. A contract designed in this way would allow both the shipper and the carrier to benefit from price changes in the marketplace while maintaining their operational relationship and associated efficiencies.

Recent evidence has suggested that shippers should hold procurement events and allow carriers to bid frequently (Martens & Suzuki, 2012). The analysis suggests holding events annually maximizes savings as it allows carriers to readjust their pricing to fit continuously changing internal and external situations and lets shippers select the most inexpensive carriers on each route. Unfortunately, holding procurement events can be expensive and time consuming. Including a spot-market index component to contracts will allow pricing to flex with market conditions and reduce some of the incentives for holding a new procurement event. The shipper would, therefore, benefit from reduced overhead costs of managing their transportation function.

2.2.4 Evaluation of contract agreements and the costs of risk-avoidance

Accurate indexes of spot-market conditions should drive better understanding of the value of contract relationships. As much as 90% of the for-hire truckload transportation market is transacted via contracts and the goal of these contracts is to both secure rates and establish longer-term relationships with carriers that will generate operational efficiencies. The alternative to these contracts is participation in the spot-market. An appropriate and accurate index of spot-market prices would provide insight for carriers and shippers about how much (if anything) each party is paying for the security of the contract.

There is an elementary relationship between contracts of different lengths. This is called the Expectation Hypothesis of Term Structure relationship and it describes that the discounted cash flows of a long-term contract should equal the discounted cash flow of a series of shorter-term contracts plus an adjustment for risk. In a simple example adapted from Kavussanos & Visvikis (2006), if a vehicle is required each and every day by a shipper, the carrier may offer a single n-period contract (e.g. 365 days). This should be equal to the sum of 365 discounted daily spot contracts plus an adjustment for risk. Mathematically:

$$\sum_{i=0}^{k-1} \frac{TC_t^n}{(1+r)^i} = \sum_{i=0}^{k-1} \frac{E_t(FR_{t+i}^m)}{(1+r)^i} + E_t\emptyset_t \quad ; \quad k = \frac{n}{m} \tag{2.1}$$

Where;

TC_t^n is the n period earnings of a contract service at time t;

$E_t(FR_{t+i}^m)$ is the expected earnings of spot charters over m periods and;

$E_t\emptyset_t$ is a term that measures risk.

k is a positive integer representing the number of spot contracts that would occur during the life of a charter contract. In this example, k would be 365. Finally, r is a discount rate applied to account for the time-value of money.

The $E_t \emptyset_t$ term can be considered as representative of the likelihood that the carrier is unable to capture a single-day contract on any given day. In a standard charter situation, the long-term contract has no risk, and therefore, $E_t \emptyset_t$ will be negative. In the case of truckload transportation, although price is guaranteed by the contract, the carrier will not necessarily receive the expected number of loads for the year. Therefore, there is execution risk on both sides of the equation and the $E_t \emptyset_t$. In this situation, the sign of the $E_t \emptyset_t$ term can be positive or negative. An alternative description of the equation might include two separate riskiness terms to make the differences more explicit.

With an appropriate index that measures spot rates over the course of a year, a shipper would be able to accurately evaluate the savings (or extra costs) of contracting prices with their carriers. Both parties would be able to determine the relative value of spot market vs. contract exposure and enter agreements that maximize their expected profitability while minimizing risk.

For more simplified negotiations, the spot market index could be used to anchor expectations in negotiations and allow companies to benchmark their contracts vs. the spot-market.

2.3 Chapter summary

This chapter developed a framework of characteristics that are important for indexes. It also reviewed the potential uses of a spot-market truckload index. In the next chapter this framework is used to review existing indexes that measure aspects of the truckload transportation market.

3. Review of existing truckload industry indexes

Several organizations offer information about truckload rates in the United States. These organizations have various stated goals and different levels of sophistication. Some explicitly state their ambition to offer an index, while others provide rate benchmarking. This section will review the information sources available within the context of the characteristics of successful indexes established in section 2.2 and listed below.

Index calculation and data collection:

1. **Accuracy:** Should be an accurate reflection of the real spot market
2. **Bias:** Should be rigorously computed and unbiased
3. **Familiarity:** Should be expressed in units familiar to the industry
4. **Balanced input:** Should be based on a sufficiently broad and balanced input
5. **Transparency:** Should be transparent and simple
6. **Disaggregation:** Should provide different levels of aggregated information in a clear and calculable hierarchy

Index management:

7. **Frequent publication:** Should be published regularly and frequently (preferably daily via electronic distribution)
8. **Auditing:** Should be audited and monitored by an independent body
9. **Complaints:** Should have proper procedures for dealing with complaints
10. **Cost:** Should be low-cost
11. **Participation:** Should be supported by the major participants in the market
12. **Updating:** Should have procedures for updating and adjusting components or index structure as market conditions change

The objective is to determine which, if any, existing indexes are appropriate for use by the industry. This section will also identify if the index discussed focuses on contract or spot-market rates.

3.1 Indexes and other market information

Available rate information can be segmented into two categories: Indexes explicitly intended to capture changes in rates over time and alternative information services. Of the indexes designed to reflect rate changes, three are well-known and widely published: The Cass Truckload Linehaul Index™, the DAT indexes, and the Stephens Freight Index. Cass Information Services (in association with Avondale Partners) publishes the Cass Truckload Linehaul Index™ based on data provided to them by large shippers. DAT publishes a number of indexes and rate information focused on the spot-market primarily derived from their experience as an intermediary in the spot-market. Stephens Research publishes a rate index based on data from large, publicly traded trucking companies. The key characteristics of these indexes are summarized in Table 3.1. Section 3.1.1 through Section 3.1.3 describes each of these indexes in more detail.

Table 3.1: Truckload rate indexes – key characteristics

	DAT	Cass	Stephens
Contract/Spot	Separate contract and spot indexes	Single index based on >90% contract	Single index based on >90% contract
Level of reporting	Separate indexes at national, regional, city and lane levels	Single index at national level	Single index at national level
Publication	Some indexes weekly, others monthly	Monthly	Quarterly
Unit of rate measurement	\$/mile	\$/mile	\$/mile
Data sources	Spot transactions, brokerage data	Shipper payments	Large carrier public data

Two other services approach rate information in entirely different ways. Chainalytics, a consultancy and information provider, maintains a rate benchmarking service where users can input the particular characteristics of a load and a benchmark rate is produced. Data is supplied primarily by large shippers. Morgan Stanley's industrial research department produces an index that tracks perceived demand and capacity by surveying participants in the industry. They do not address rates directly, but instead, imply that rate trends can be inferred by changes in industry demand and capacity. Section 3.1.4 and 3.1.5 discuss each of these services in more detail.

3.1.1 DAT Trendlines and Rateview™

DAT began as "Dial-a-Truck" in 1978 as a load finder service. Over the years, technology has advanced and DAT now manages an online loadboard on which more than 68 million loads are offered each year. DAT offers rate information for both contract and spot rates, but maintains separate indexes for each. They offer two major products related to spot-market rates.

DAT Trendlines provides information on changes in regional and national rates on a weekly basis. This product highlights particular lanes that are outliers from the norm within their established five regions. DAT explicitly segments its rates by equipment type and also separates contract rates from spot-market rates. The information is expressed in rate per mile. DAT Trendlines is widely available at no charge.

DAT offers, for a fee, a more comprehensive product with DAT RateView™. RateView is based on data from a number of contributors, including freight brokerages. DAT separates the contiguous United States into 133 markets and provides more complete and comprehensive information about specific origin-destination combinations (e.g. Boston region to Los Angeles region) where there is sufficient data available. They also display the number of datapoints used to arrive at the rates. This is important for users to gauge the accuracy of the rate calculation. RateView also provides access to rate information over different time scales – from the most recent week up to the most recent year.

The combined services offered by DAT address many of the key characteristics of successful indexes. They provide several different levels of aggregation that lets the user determine the information that best works for their business. DAT is an independent organization and the basic information provided is free. Rate trends are updated weekly, but more current information is available by subscription through RateView™. Brokerages and other contributors earn credits to access rate information by providing their transaction data.

While it is clear how rates are calculated on individual lanes, there is, perhaps, some need for greater clarity in DAT's descriptions of how the aggregated indexes are calculated. There is also a need for DAT to ensure that its data sources are comprehensive and representative of the industry at large. DAT also appears to lack formal input and participation from the industry in the design and ongoing management of its information services. While this does not suggest that DAT is not managing the index well, it is a significant difference from the indexes that have developed successfully in other industries. Figure 3.1 shows a recent map by DAT of prevailing outbound rates by region. Figure 3.2 shows nationally aggregated van rates per mile over the past year with comparison to the contract market.

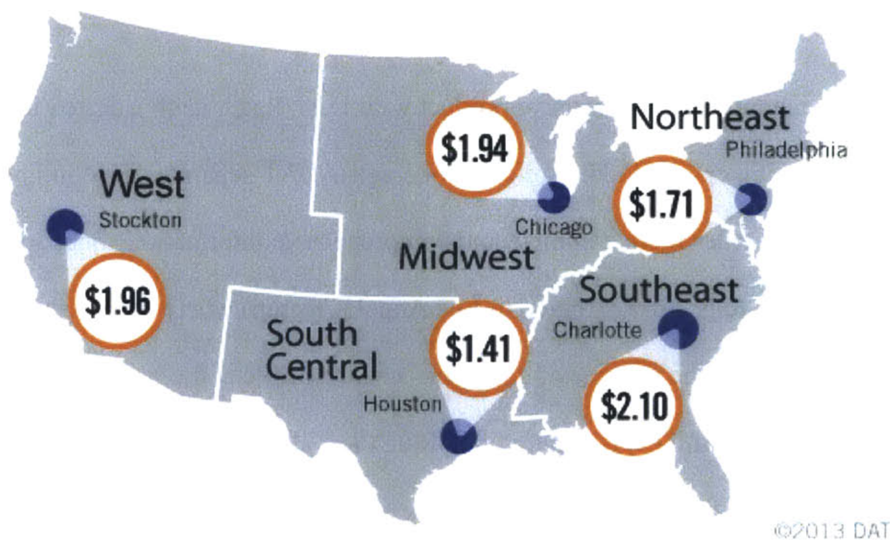

Figure 3.1: DAT Rate per mile index for each region, March 2013

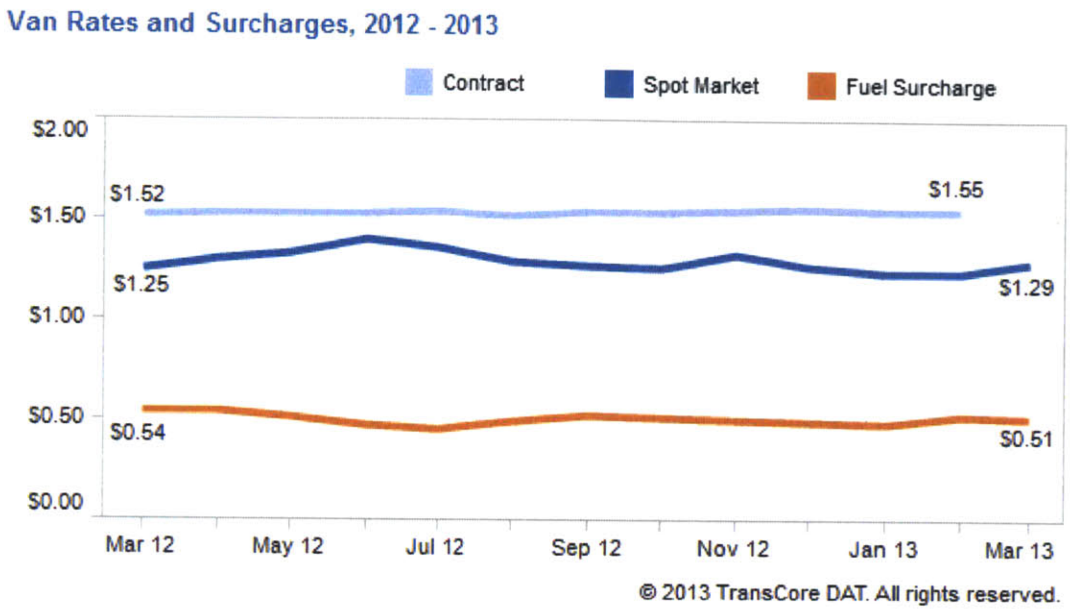

Figure 3.2: DAT national rate per mile index – 2012-2013

3.1.2 Cass Truckload Linehaul Index™

Cass Information Systems provides payment services to the trucking industry and, therefore, has access to massive data sets. The Cass Truckload Linehaul Index™ is calculated from more than $17 billion in payments processed annually. The index is a measure of per-mile rates excluding fuel and accessorials for dry van truckloads. Cass excludes fuel and accessorials based on the data recorded on each transaction they receive from their customers. They do not adjust for different fuel surcharge programs.

The Cass index has a number of attractive qualities. The dataset is large and spread across a range of shippers from various industries and geographic locations. However, the bulk of their transactions are concentrated within larger Fortune 500 companies. Spot-market transactions are included in their analysis, but represent only 5-10% of all transactions. No separate rate index is published that focuses exclusively on spot rates. The index itself uses a base of 100 in 2005 and changes are measured in relationship to that base. Cass publishes the information monthly and the data is quite current. Their stated goal is to publish the index within 3-5 days of the end of each monthly period.

However, the index has a number of limitations. Using the base unit of rate per mile can lead to distortions if the average length of haul is not constant. Over the last few years, the industry has seen greater competition with intermodal providers on longer routes and this has reduced the average length of haul in trucking. There is a commonly understood inverse relationship between length of haul and rate per mile that can be explained by the higher efficiencies and less stop-time involved in longer hauls. Additionally, with shorter average haul lengths, a carrier's fixed costs must be absorbed by fewer miles. Therefore, as the average length of haul in the industry shrinks, there will be an expected increase in the average rate per mile. This change is independent of any changes in rates on any individual route. A shipper with a network that has not changed in a similar way will not experience the change in rates that are the result of this effect in the index.

For the purposes of an index that tracks the spot-market, the dataset used, the monthly publishing, and the aggregation of data are limiting. Cass collects payment information from many organizations, but most of the data is based on contracted rates. Monthly information is useful for strategic decision-making, but not the more tactical, hands-on uses discussed in Section 1. Finally, the Cass index provides a high-level view of national rates. It cannot be used for determining rates on specific lanes or to compare changes in rates between regions. Most shippers and carriers manage a specialized network of routes. They are exposed to local and regional effects as well as national changes. Because Cass does not provide sufficiently disaggregated data, the experiences of shippers and carriers will not likely match changes in the Index. Figure 3.3 shows the history of the Cass Truckload Linehaul Index™ since inception.

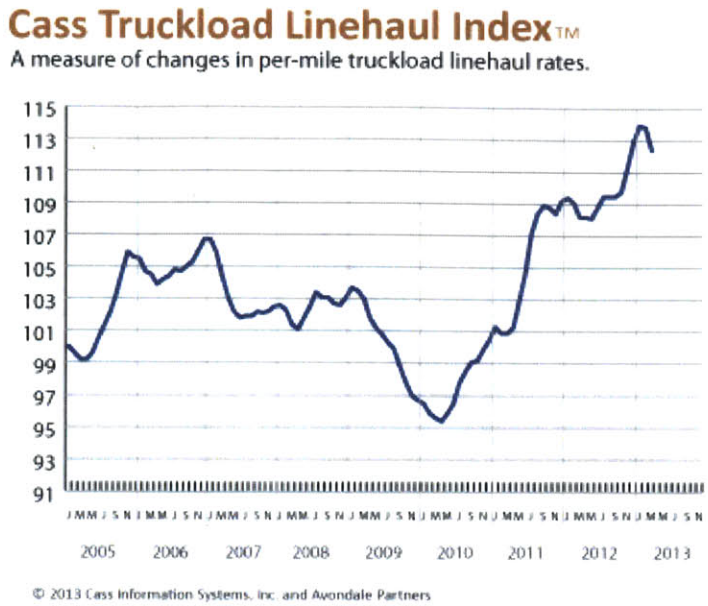

Figure 3.3: Cass Truckload Linehaul Index™ – 2005-2013

3.1.3 Stephens Freight Index

The Stephens truckload index focuses on the largest, publicly traded carriers. It uses available information to calculate the average rate received per mile by the large carriers and provides quarterly updates. While this approach provides useful insight into the profitability and operations of large carriers, it has less applicability for the broader spot market.

Similar to the Cass Index, the underlying unit is rate per mile and the Stephens index explicitly acknowledges that its results can be distorted by changes to the average length of haul in the industry. Between 2006 and 2012, the average length of haul in their dataset dropped from 759 miles to 611 miles. Regardless of the changes in rates on any particular route, the change in average length of haul would show up as an increase in rate per mile. The structure of the Stephens Index makes it challenging to disentangle this effect from real rate changes. Figure 3.4 shows the Stephens truckload index since inception. The established base of the index in 1995 was 100 units.

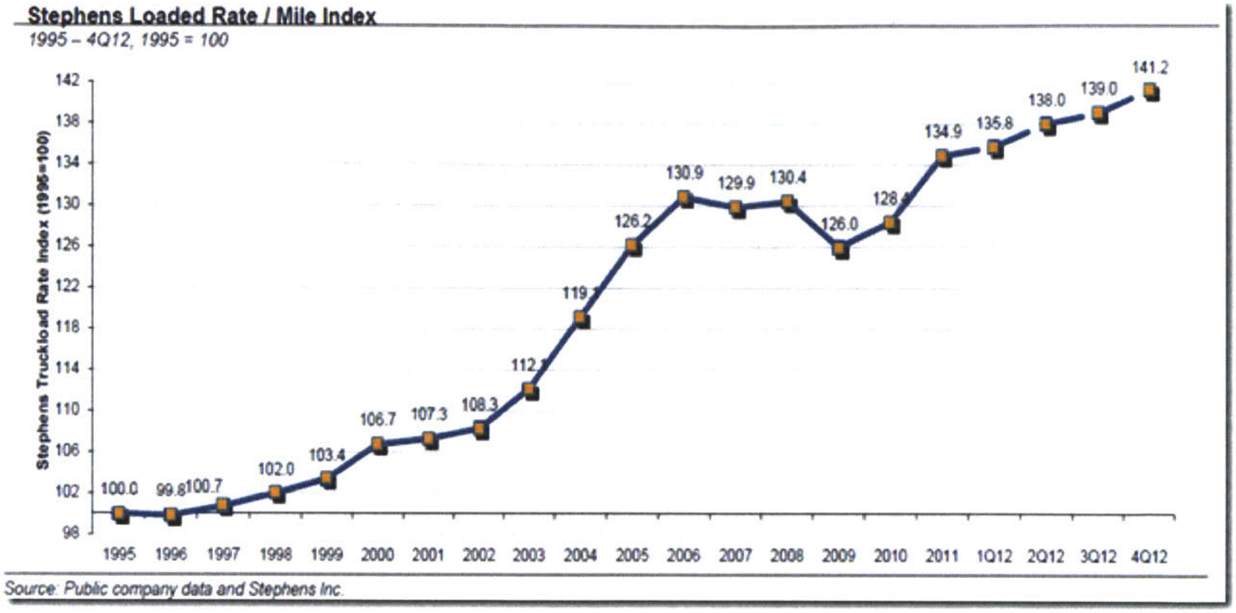

Figure 3.4: Stephens Freight Index – 1995-2012

3.1.4 Chainalytics

Chainalytics manages the Freight Market Intelligence Consortium, a group of shippers that provide information on the rates they pay on thousands of different lanes across the country. They collect the information, normalize the data to remove accessorial and other unusual charges and then analyze the impact of various load and company characteristics. These characteristics include origin and destination locations, length of haul, equipment characteristics, spot vs. contract carrier moves, the shipper's annual volume in that lane, the fuel surcharge programs involved, the speed of service and the presence of dangerous materials in the load. Chainalytics creates an econometric model of the market based on the coefficients of the analyzed characteristics. When a user of the benchmarking service wants to understand rates in a lane, they input their own particular load characteristics and the Chainalytics model arrives at a benchmark value.

The Chainalytics data is sufficiently disaggregated to be useful to companies with different route networks. The benchmarking model is updated monthly, but it includes data from prior periods, so discerning seasonal spot market changes is difficult. Data is sourced from a diverse set of member companies from many industries, but is focused on contract market rates.

Separately from the benchmarking model, Chainalytics provides weekly updates on rate trends specified by equipment type and industry. Chainalytics charges a membership fee for participation in the consortium. This poses a barrier to its widespread adoption and use. Chainalytics' approach to normalization of data can provide good insights, however, into how an index for the spot market might be calculated.

3.1.5 Morgan Stanley Freight Indexes

Morgan Stanley indexes use an entirely different approach to understanding changes in rates. They do not examine rates directly; instead they focus on changes in capacity and demand for each of the common equipment types. With this information, a user can then infer the future direction of rates. Capacity and demand values are determined via simple surveys collected from a variety of people familiar with the industry. Because these indexes do not provide clear information on rates, they are not appropriate for the uses discussed in Section 2.2. However, in a market with functioning freight forward agreements and futures, the information from Morgan Stanley would be useful for industry users and speculators as they examine the potential for future rate changes. Figure 3.5 shows an example of the index for dry vans from 2007 to 2012. Of note is the very low level of incremental demand/incremental supply during the 2009 recession.

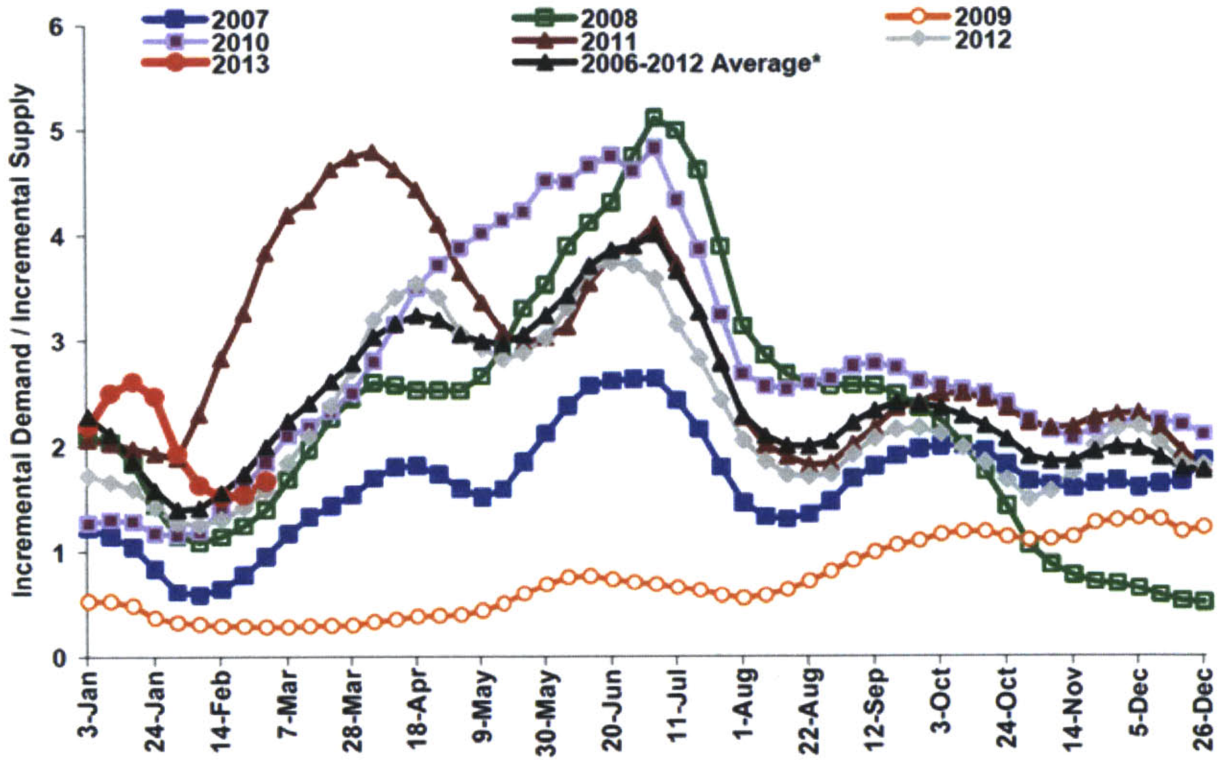

Figure 3.5: Morgan Stanley Freight Index 2007-2013

3.2 Comparison of index characteristics

Table 3.2 summarizes the characteristics of the discussed indexes within the framework discussed in Chapter 2. Overall, the indexes are stronger in the index calculation and data collection than in the index management half of the framework. Although they are not indexes, the Chainalytics and Morgan Stanley products are included in the framework for comparison purposes.

Table 3.2: Existing index comparison

	Characteristic	Indexes			Other	
		Cass	DAT	Stephens	Chain-alytics	Morgan Stanley
Index calculation and data collection	1) Accuracy: Should be an accurate reflection of the spot market		✓			
	2) Bias: Should be rigorously computed and unbiased		✓		✓	
	3) Familiarity: Should be expressed in units familiar to the industry	✓	✓	✓	✓	
	4) Balanced input: Should be based on a sufficiently broad and balanced input	✓	✓		✓	✓
	5) Transparency: Should be transparent and simple	✓	✓	✓		
	6) Disaggregation: Should provide different levels of aggregated information in a clear and calculable hierarchy		✓			
Index management	7) Frequent publication: Should be published regularly and frequently (preferably daily via electronic distribution)		✓			
	8) Auditing: Should be audited and monitored by an independent body					
	9) Complaints: Should have proper procedures for dealing with complaints					
	10) Cost: Should be low-cost	✓		✓		✓
	11) Participation: Should be supported by the major participants in the spot market					
	12) Updating: Should have procedures for updating and adjusting components or structure as market conditions change				✓	

Of the three major indexes, DAT aligns best with the characteristics of freight indexes that were discussed in Chapter 2. They are the only organization to develop separate spot market indexes that do not include contract rate data. Similar to the Baltic Indexes, they provide information at different levels of aggregation including rate reports at the lane level, the regional level and the national level. The Cass Truckload Linehaul Index™ and the Stephens Freight Index are not appropriately designed for use in the spot-market. They sample primarily contract rate transactions and provide only a single, national level rate index.

The clear lesson from experiences in other industries is the need to segment data that has different characteristics. Chainalytics does this by examining the impact of each characteristic on the cost of transportation and then estimating a new cost based on inputted values. It, however, does not show how the impact of each characteristic changes over time.

On the index management side of the framework, it is clear that no existing index shares the characteristics that have proven successful for the Baltic Indexes and indexes in other industries. This may be attributable to the different circumstances that were involved in the development of those indexes versus the development of truckload indexes in the United States. The Baltic Indexes developed through organized, concerted effort by major industry players. They established clear governance procedures and were concerned about ensuring the independence of the index. Because the major industry players were involved in the design of the index, they have continued to support its growth and development. A similar process of industry-supported index development has not occurred in the truckload spot market in the United States. The truckload indexes discussed in this chapter have been developed primarily by independent analysts and for-profit enterprises.

3.3 Chapter summary

This chapter reviewed several transportation indexes that are currently published in the United States. It examined their characteristics in comparison with the framework of successful indexes developed in Chapter 2. Of the examined indexes, the DAT freight indexes exhibited the most characteristics in common with spot rate indexes in other industries. No index exhibited all of the characteristics of the framework and no index demonstrated the index management structure that has been successfully used by the Baltic Index.

4. Profiling the dataset

Throughout this analysis, data has been provided by LogisticsCo, a large North American provider of non-asset based logistics services. Within their business they have identified a need for an index of spot-market rates and have noticed that existing indexes in the market have failed to capture some of the trends and dynamics they have experienced. To help analyze this problem, the company has provided a dataset of nearly one million movements from the spot-market portion of their business from October 2009 until October 2012.

4.1 Data preparation

Several steps were taken to ensure that the data used would be as accurate and useful as possible. Each transaction nominally represents an actual truckload movement; however, occasionally transactions are used for adjustments or refunds. These types of transactions were excluded from the analysis. Movements with destinations or origins outside the contiguous United States were also excluded. Furthermore, all equipment types with the exception of Dry Vans and Reefers were excluded as none had sufficient data to make analysis effective. Altogether these exclusions accounted for approximately 7% of the original dataset.

After examining portions of the data, it is clear that strong seasonality characteristics exist in some regions. To prevent a particular season from having greater weighting in the analysis than other seasons, all data analysis was restricted to the period October 1, 2010 to September 30, 2012. This two-year study period also allows for visualization of significant repeating seasonal variability.

Throughout the study period, LogisticsCo has been expanding their operations. Average monthly transaction volume in the dataset at the end of the study period is approximately twice that at the beginning of the study period. This growth may create challenges in the analysis due to changes in customer mix and volume that impacts rates.

4.2 Data description

Each transaction includes a number of fields that align with characteristics that drive truckload transportation costs. These include trip miles, equipment type, and fuel price. A number of different origin and destination coding is available in the data. Specific longitude, latitude and location data is available, but each transaction is also coded with an "origin code" and "destination code" that corresponds to a metropolitan area. The dataset has 127 of these codes. Figure 4.1 shows the Chicago, IL Origin Code region. The relative size of the circles is indicative of the number of loads that originate at each location. With 101 miles between the two furthest locations, this region is relatively compact.

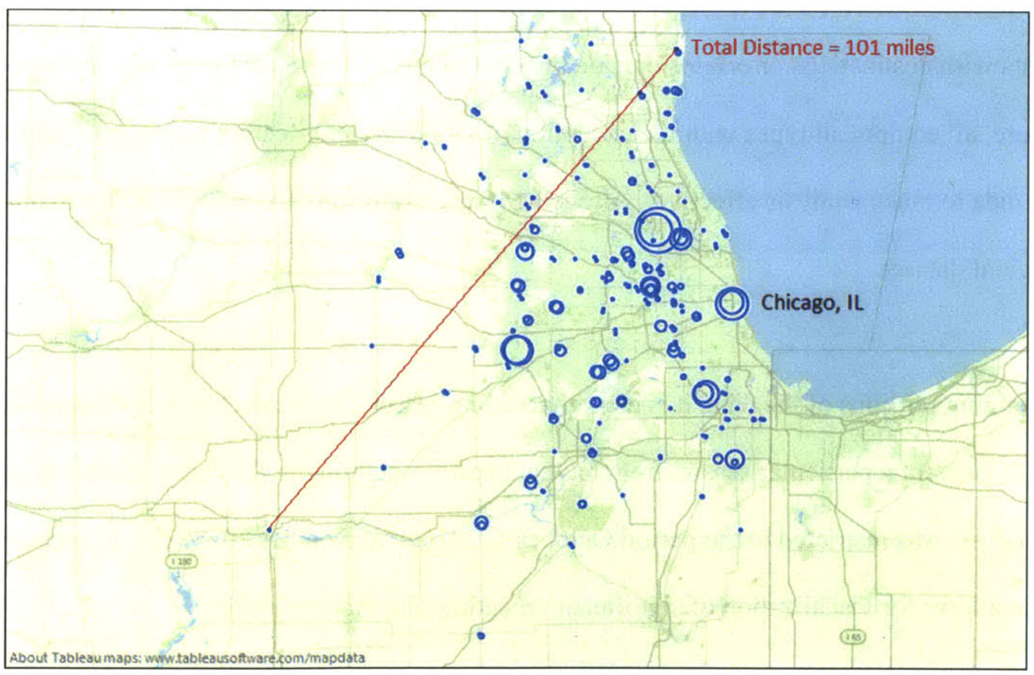

Figure 4.1: IL-CHI Origin Code

Figure 4.2 shows the Portland, OR Origin Code region. This region is larger and spread along 248 miles of the Willamette and Columbia River valleys. A combination of origin code and destination code will be described as a corridor for the remainder of this paper.

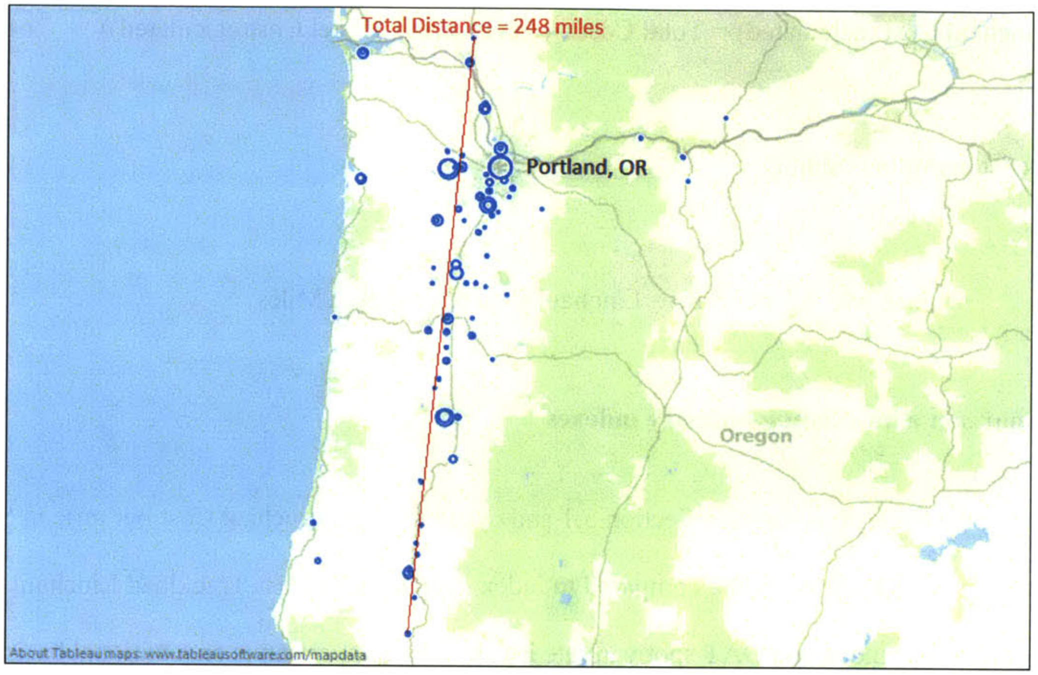

Figure 4.2: OR-POR Origin Code

The dataset explicitly indicates a "Total Cost". It also has a separate field that indicates the value of all accessorials that were charged in each transaction. Some transactions also explicitly indicate the fuel component of total cost. However, a significant portion of the transactions used a flat rate that did not break out the fuel component. To understand the line-haul cost (the portion of cost that is not accessorial or fuel related), this had to be normalized. A calculation of fuel cost was created. It calculates the expected fuel cost by dividing the Miles recorded in the transaction data by the Federal Highway Administrations estimate of fuel consumption for combination trucks (5.9MPG) and then multiplying by the prevailing National Diesel rate (also recorded in the transaction data). The equation for the fuel cost calculation is shown below.

$$\text{Fuel Cost (calculated)} = \text{Miles}/5.9\text{MPG} * \text{Weekly National Diesel rate} \quad (4.1)$$

Linehaul cost was then calculated by subtracting the Accessorials cost available in the dataset and the fuel cost calculated in equation 4.1 from the Total Cost value from the dataset:

$$\text{Linehaul Cost (calculated)} = \text{Total Cost} - (\text{Accessorials} + \text{Fuel Cost(calculated)}) \qquad (4.2)$$

Linehaul cost per mile is simply:

$$\text{Linehaul Cost per mile} = \text{Linehaul Cost (calculated)/Miles} \qquad (4.3)$$

4.3 Comparison with existing industry indexes

After making the exclusions noted in Section 5.1 and calculating the Linehaul Cost per mile in Section 5.2, the rates from LogisticsCo were compared to index data from the Cass Truckload Linehaul Index™, the Stephens Freight Index and DAT spot van rate index. The comparison is on a national basis. An October 2010 baseline of 100 is used and the results are shown in Figure 4.3.

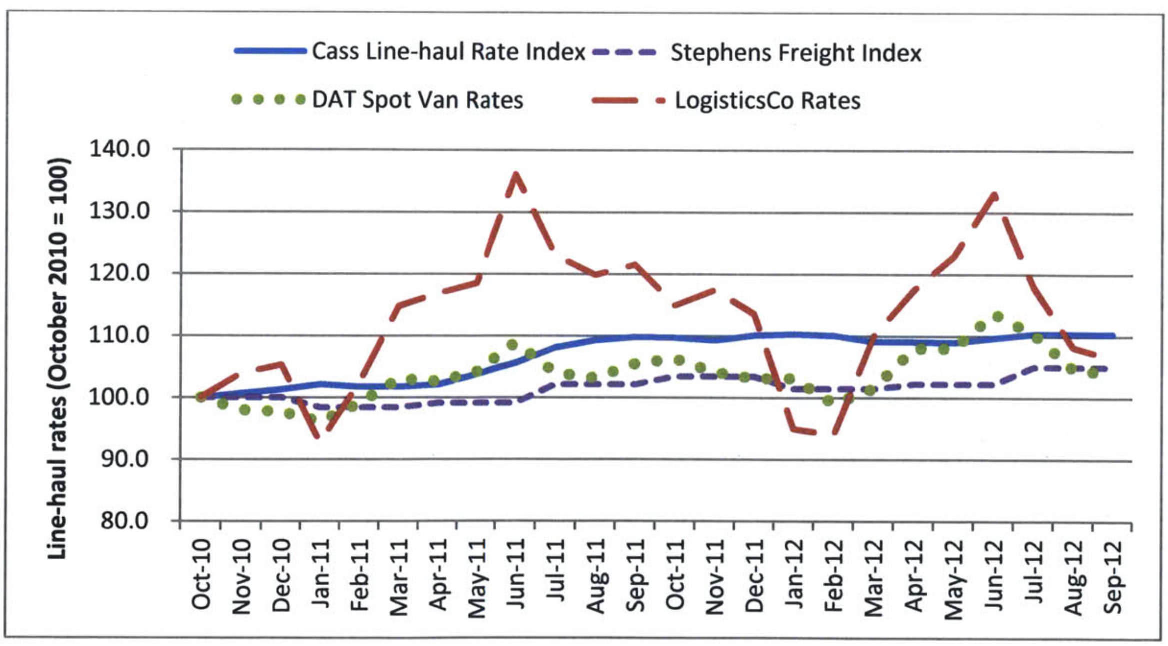

Figure 4.3 Comparison of LogisticsCo data with other indexes

The Cass Truckload Linehaul Index™ and the Stephens Freight Index are very stable and do not exhibit noticeable seasonality. This is likely the result of their datasets including contract rates. By contrast, both DAT and the LogisticsCo data show seasonal peaks in June each year and seasonal troughs in January and February. The scale of the seasonal impact, however, differs significantly. DAT shows a range from peak to trough of approximately 14 index points while the range for LogisticsCo is 41 points.

It appears that the LogisticsCo data measures much more significant seasonal price changes than are measured by DAT. The DAT index may be more stable as it is based on data from a number freight brokers and includes roughly twenty times the number of transactions as LogisticsCo in each period. The extra seasonality in the LogisticsCo data may also indicate that their business is more concentrated in areas with seasonal rate patterns. It might also indicate that some other characteristic of their business operations makes their rates more sensitive to changing conditions. The remainder of this chapter describes two primary candidates that may explain the discrepancy between the rates experienced by LogisticsCo and the rates measured by DAT. Section 2.4 examines the geographic distribution in the dataset. Section 2.5 examines the distribution of lead-times.

4.4 Geographic distribution

With 127 origin/destination codes, there are a possible 16,129 corridors. During the study period, 8434 of those corridors (52.3%) experienced volume. The maximum number of transactions for any single corridor (Orlando, Florida to Miami, Florida) was 4,564 over two years or approximately 9 per day. The distribution of transactions to corridors is quite concentrated. 3.4% of the corridors account for 50% of transactions and 26.7% account for 90% of all transactions. The distribution of total dollars spent is slightly less concentrated. 5.2% of the corridors account for 50% of total spend and 33.3% of corridors account for 90% of total spend. Figure 4.4 shows the cumulative distribution of dollars spent.

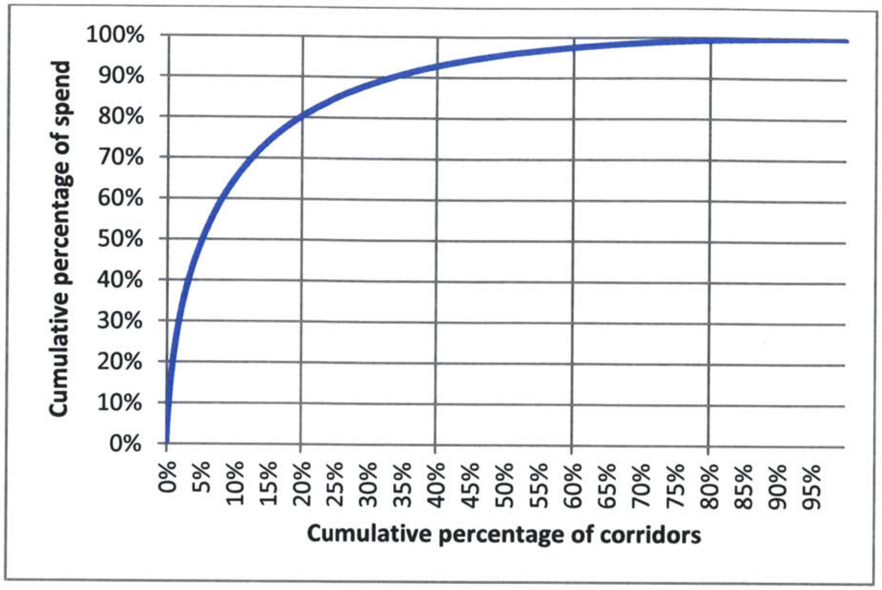

Figure 4.4: Cumulative distribution of total dollars spent by corridor

Figure 4.5 shows the distribution of load origins by state. While LogisticsCo had transactions in all 48 contiguous states, they have high concentrations in Texas, Georgia, Florida and California. Darker colors indicate states with higher volume and the numbers indicate the percentage of total volume that originates in each state.

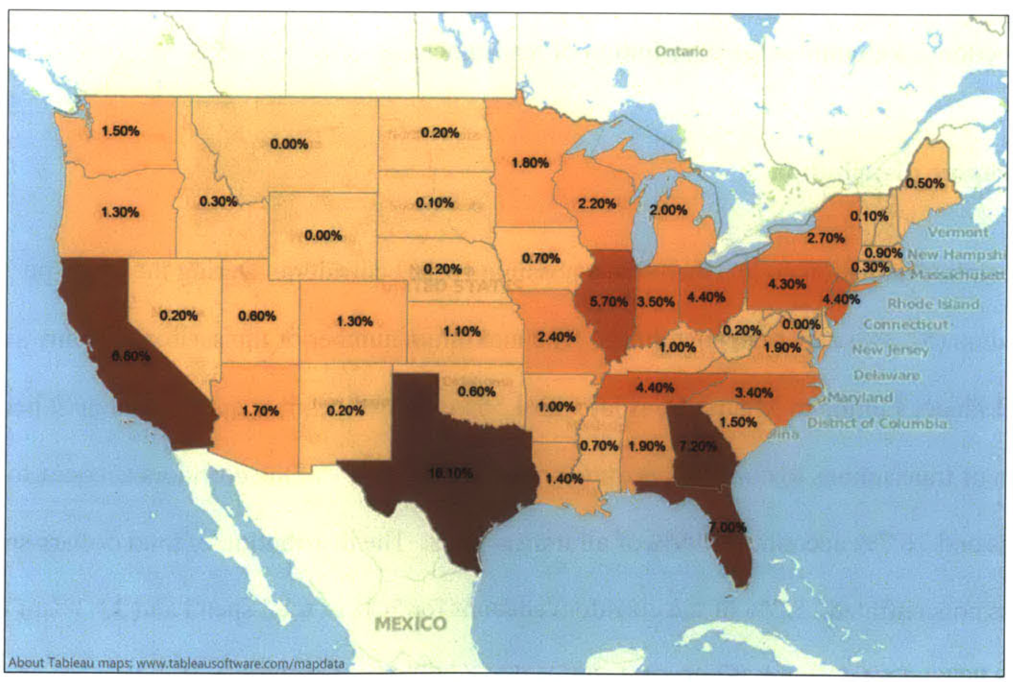

Figure 4.5: Distribution of load origins by state

Figure 4.6 shows the distribution of load destinations. LogisticsCo has proportionally fewer destinations in Texas and Georgia, and more loads destined for the Midwest states of Illinois, Indiana and Ohio. Again, darker colors indicate states with higher volume and the numbers indicate the percentage of total volume that originates in each state.

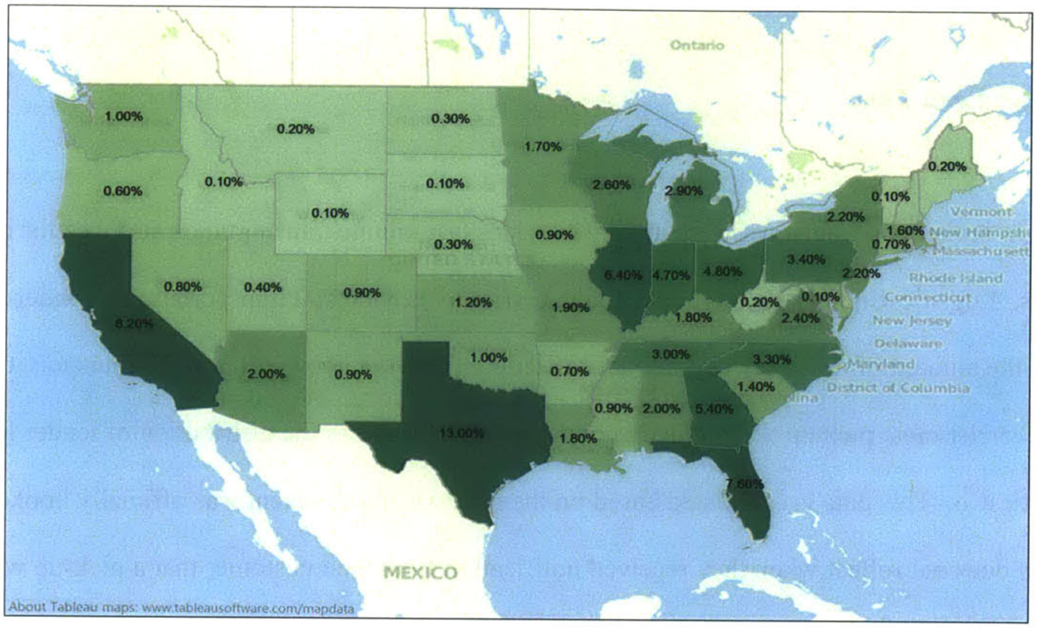

Figure 4.6: Distribution of load destinations by state

The geographic distribution of loads shows that LogisticsCo has significant exposure to traditional agricultural regions like California, Texas and the Southeast that experience constrained capacity during harvest seasons. The distribution also provides evidence that LogisticsCo derives a greater proportion of their revenue from these regions than does a major competitor. Abramson & Sawant (2012) used a dataset of 600,000 loads from a competitive freight brokerage and split the country into six regions. They determined that 19% of loads originated in the South-central region (Colorado, New Mexico, Texas, Oklahoma, Kansas, Missouri, Arkansas and Louisiana). That same region provides 23.5% of the loads in the LogisticsCo dataset. Abramson and Sawant's Northeast region contains all states east of Illinois and north of North Carolina and was the origin of 32% of the competitor's loads. For LogisticsCo, only 26% of loads originated in this region.

LogisticsCo derives proportionally more of their business from agricultural regions and proportionally less from the Northeast than one of their major competitors. This may partially explain the greater seasonal rate changes present in the LogisticsCo data.

4.5 Tender Lead Time

As a freight brokerage, LogisticsCo frequently provides last-minute transportation services for its customers. LogisticsCo must then work to find a carrier to provide that transportation. Tender lead time describes the amount of time between when a carrier confirms that they will provide transportation and when the carrier must pick up the shipper's goods. Figure 4.7 shows the distribution of tender lead times for LogisticsCo. This data is calculated based on the time when a shipment was officially booked in their system. It does not reflect when they received notification from their customer that a pick-up was needed.

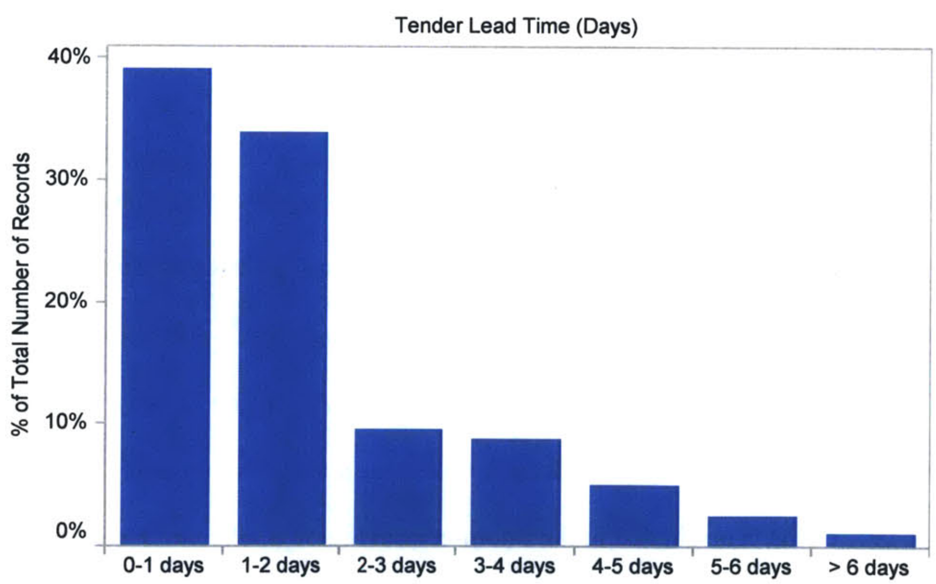

Figure 4.7: Distribution of tender lead time in dataset

This distribution differs markedly from the tender lead time distribution shown in previous studies. Abramson & Sawant (2012) found that a competitive freight brokerage had 28% of transactions booked on the same day as pick-up and 33% of loads had lead times of more than three days. In LogisticsCo's dataset, 39% of loads were same-day and only 17% of loads had lead times of more than three days. Caldwell & Fisher (2008) examined a different set of data representing approximately 330,000 from another freight brokerage. They found an average tender lead time of between 2.5 and 3 days. The comparable figure at LogisticsCo is 1.7 days.

LogisticsCo appears to operate with shorter lead times than a major competitor. If loads with short lead times are associated with a greater increase in rates during seasonal peaks than are loads with longer lead times, LogisticsCo's average rates would experience more seasonal rate volatility than other brokers. This could further explain the discrepancy between LogisticsCo's seasonal rate patterns and the rate patterns measured by DAT.

4.6 Chapter summary

In this chapter, the source of the data used for this analysis and the steps that were taken to develop the final dataset were described. Existing indexes were compared to a simple average rate index developed from the LogisticsCo dataset. The data from LogisticsCo experienced significantly more seasonal volatility than existing indexes. The geographic distribution of LogisticsCo's business and their short tender lead times were discussed as possible explanations. The next chapter will examine in more detail the relationship between geography and rates.

5. Geographic disaggregation - Corridors

One of the key characteristics of indexes defined in Chapter 2 was that an index should provide different levels of aggregated information. In particular, indexes from other industries often provide information disaggregated by geography. For truckload transportation, the corridor is a convenient level of disaggregation. However, the necessity of providing information on rates for each corridor is contingent on different corridors having different rate behaviors. If all corridors behave similarly, a single national index would be sufficient.

This chapter reviews the rationale for using corridors to disaggregate by geography. It also explains why line-haul rate per mile is appropriate for measuring rates in a corridor. Finally, it examines the rate behaviors of several corridors to determine if this level of disaggregation is appropriate.

5.1 Corridors as the basic level of disaggregation

The corridor is a convenient basic unit for geographic disaggregation for several reasons. The LogisticsCo dataset has 8,434 corridors with volume. Of those, only 410 have volumes of five or more loads per week. Corridors with less volume than that would prove challenging for developing a frequently published index. Work in the contract market has indicated that further disaggregation to the lane level (typically a combination of a 5-digit zipcode origin and a 5-digit zipcode destination) is useful. However, disaggregating to this level would reduce the number of lanes that have sufficient volume for indexing and data at the lane level might be heavily concentrated with a single customer. Rates on a lane with high customer concentration may be related to the particular behavior or needs of the shipper and changes in rates may not be representative of overall market conditions. Calculating rates across an entire region reduces the likelihood of a strong influence from any single particular customer overwhelming the influences of changes in the overall market. Figure 5.1 summarizes the trade-off between these factors.

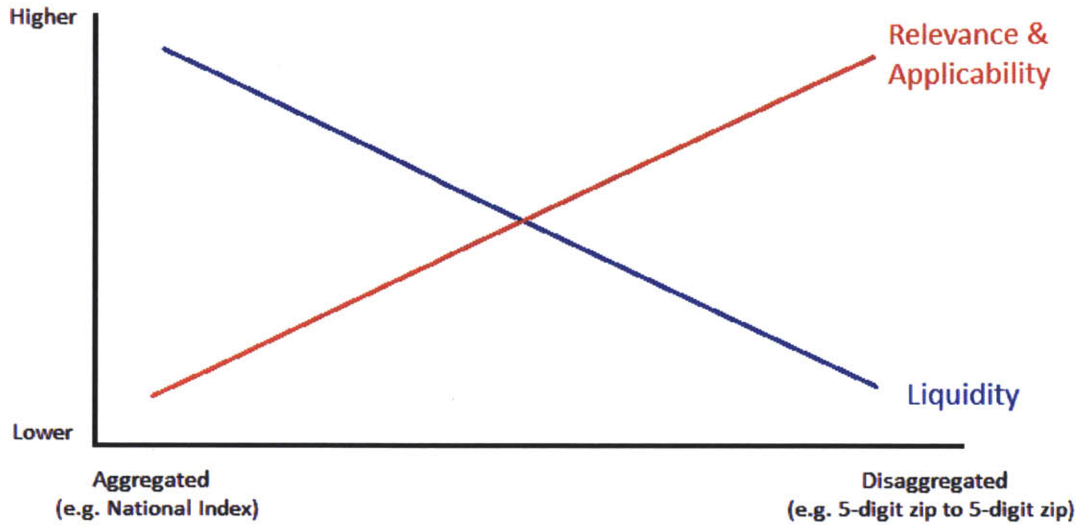

Figure 5.1: Trade-off between relevance and liquidity of disaggregation

As indexes become more disaggregated, the data becomes more relevant and meaningful to a particular user. However, using disaggregated data means generating indexes from small, extremely specified markets with little liquidity. The use of corridors is a middle-ground between national level indexes and rate indexes generated for individual 5-digit zipcode to 5-digit zipcode lanes.

5.2 Rate units

As noted in section 4.3, origin destination codes can be quite large. In the western United States, many regions extend over several hundred miles. Therefore, two transactions, both sharing the same origin and destination codes, may have significantly different length of haul. To account for this in measuring rates on a corridor, it is necessary to use units of rate per mile.

The total rate of the transaction is comprised of fuel costs, accessorial fees and the line-haul cost. It is the line-haul portion that is of most interest to the trucking industry. Accessorial fees depend on particular actions that occur on a particular load. They are fees for services performed beyond a standard pick-up and delivery. A typical accessorial fee might be for "detention" – time spent by the carrier at the shipper's location beyond a specified threshold. Fuel costs are well-known, tracked and understood by various market participants. The line-haul portion is the piece that shippers, carriers and brokers look to manage for their profitability.

Based on these considerations, the most effective unit of measurement for corridor rates is line-haul rate per mile. Per mile rates are used frequently by several of the existing truckload indexes and are easy to understand. The also allow for easier comparisons between corridors.

5.3 Examining rate behaviors by corridor

To test the behavior of different corridors, some basic characteristics of ten corridors across the United States were reviewed. The chosen corridors have high transactional volume and were chosen to represent a range of geographies and length of haul. Table 5.1 describes the basic characteristics of these corridors. Corridors have been assigned numbers and their locations are described as regions to protect confidential rate information. These corridors range in length from 264 miles to 972 miles. Of note, corridors 9 and 10 are a complementary pair. The origin of corridor 9 is the destination for corridor 10 and vice versa.

Table 5.1: Selected corridors and key characteristics

Corridor #	Region	# of Observations	Length of haul (miles)	Average Line-haul rate ($/mile)	25th percentile rate ($/mile)	50th percentile rate ($/mile)	75th percentile rate ($/mile)	Coefficient of variation
1	SE	2254	453.2	1.78	1.47	1.67	1.92	0.23
2	SW	2421	358.4	1.76	1.35	1.64	1.89	0.27
3	Midwest-SE	251	753.8	1.23	1.02	1.14	1.32	0.18
4	NE-Midwest	2049	802.4	0.38	0.27	0.34	0.42	0.27
5	S	1709	264.2	1.99	1.50	1.80	2.19	0.29
6	NW	911	790.8	1.52	1.23	1.31	1.56	0.29
7	Midwest	1780	402.2	1.20	0.91	1.05	1.25	0.31
8	E	733	516.6	2.12	1.85	2.07	2.21	0.20
9	W	1362	962.5	0.46	0.30	0.36	0.46	0.44
10	W	666	972.3	1.54	1.22	1.40	1.73	0.28

The remainder of this chapter examines four characteristics of rate behavior to establish whether the corridors are broadly similar or broadly dissimilar: overall average rates, volatility, seasonality, and year-over-year trend.

5.3.1 Differences in magnitude of weekly average rates

There is a dramatic difference in average rate between different corridors. Figure 5.2 shows the weekly average rate for all ten corridors. For the two year period, Corridor 8 has the highest average rate per mile ($2.12) while Corridor 4 has the lowest average rate per mile ($0.38). The highest average rate in a specific week was Corridor 5 in the week of June 26, 2011 when rates peaked at $2.75 per mile. The lowest average rate in a specific week was Corridor 9 in the week of June 26, 2011 when rates bottomed out at $0.30 per mile.

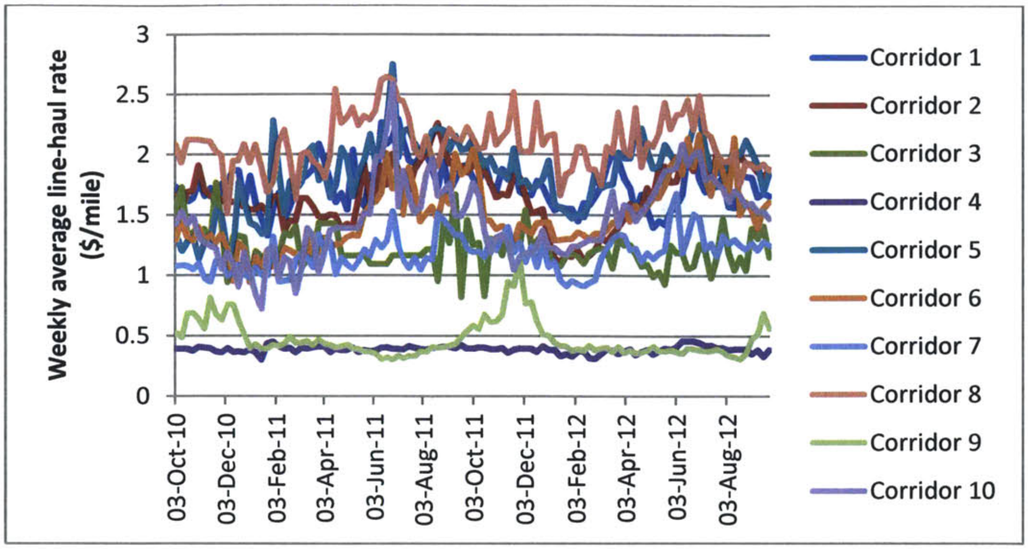

Figure 5.2: Average weekly rates for all corridors

Figure 5.3 shows a 4-week moving average for each corridor. While two corridors have single peaks in June (Corridors 8 and 10), others appear to have multiple peak periods (Corridors 1 and 6). Corridors 3, 4 and 7 have little seasonality. Corridor 8 has the highest average rates in nearly all periods while Corridors 4 and 9 consistently have the lowest rates.

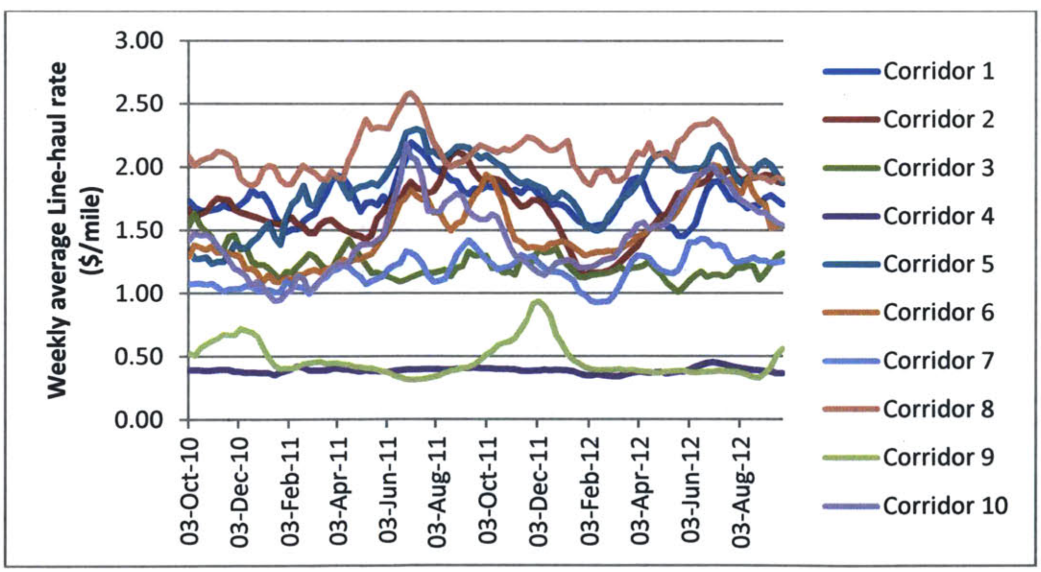

Figure 5.3: Average weekly rates for all corridors (4-week moving average)

5.3.2 Differences in variability of weekly average rates

Variability of average weekly rates was further examined for each of the corridors and a summary of these statistics are listed in Table 5.2. Average weekly change describes the average absolute percentage change in the weekly rates. Of note, most corridors had average weekly changes of between 5 and 10%. Corridor 3 is an anomaly with much higher average weekly moves. It is also the lowest volume lane in the data set and demonstrates the instability of a measure of rates when there are relatively few data points. Peak and trough rates are compared to a base level of 100 that represents the average rate over the two-year study period. Therefore, a peak rate of 131.4 means the highest weekly rate during the 2-year study period was 31.4% higher than the average rate on that corridor.

Table 5.2: Measures of variability on studied corridors

Corridor #	Region	Avg Weekly Change	Trough Rates	Peak Rates	Range
1	SE	5.7%	75.9	131.4	55.5
2	SW	4.6%	66.4	135.2	68.8
3	Midwest-SE	14.3%	66.9	144.9	78.0
4	NE-Midwest	5.8%	77.8	118.8	41.0
5	S	8.4%	72.8	149.3	76.5
6	NW	7.3%	63.5	146.2	82.7
7	Midwest	7.3%	77.4	142.9	65.5
8	E	6.8%	72.0	126.1	54.1
9	W	9.3%	63.7	239.5	175.8
10	W	9.9%	49.3	176.8	127.5
Average		**7.9%**	**68.6**	**151.1**	**82.5**

These measures reveal that volatility is dramatically different between different corridors. Corridor 9 has a range of rates that is nearly twice as large as its average rates. Figure 5.4 shows the weekly rates for all ten corridors normalized so that the average rate for the two-year period for each lane is 100.

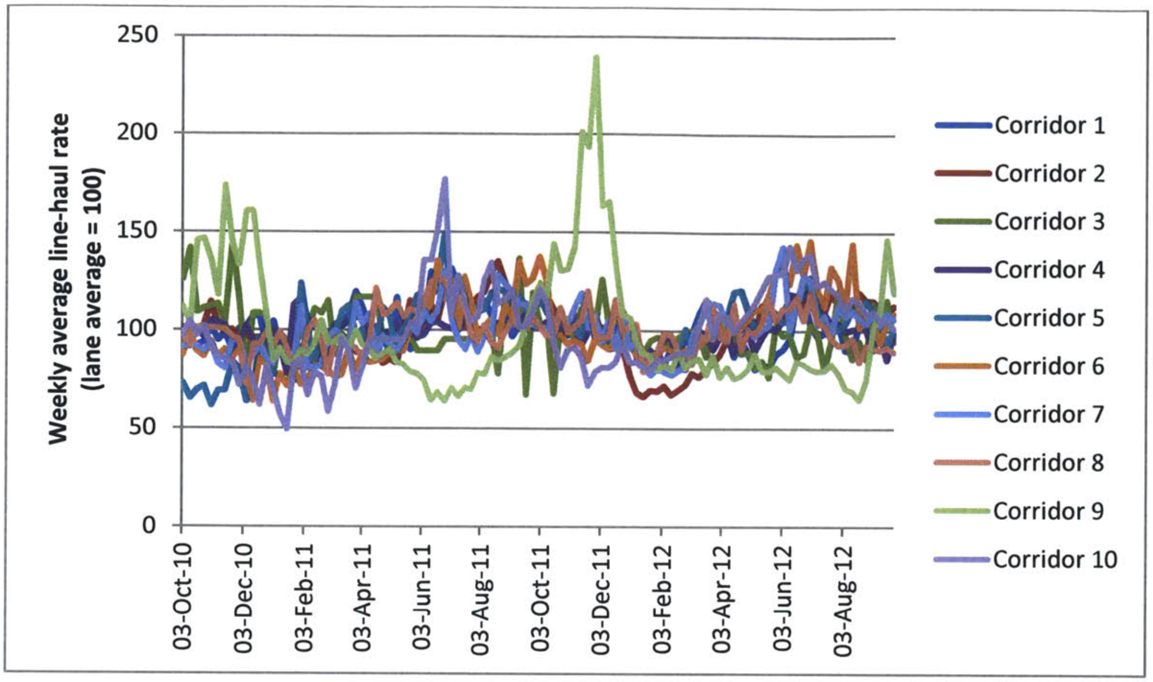

Figure 5.4: Normalized average weekly rates for all corridors

Figure 5.5 shows a comparison between the corridor with the smallest range of rates (Corridor 4) and the corridor with the largest range of rates (Corridor 9).

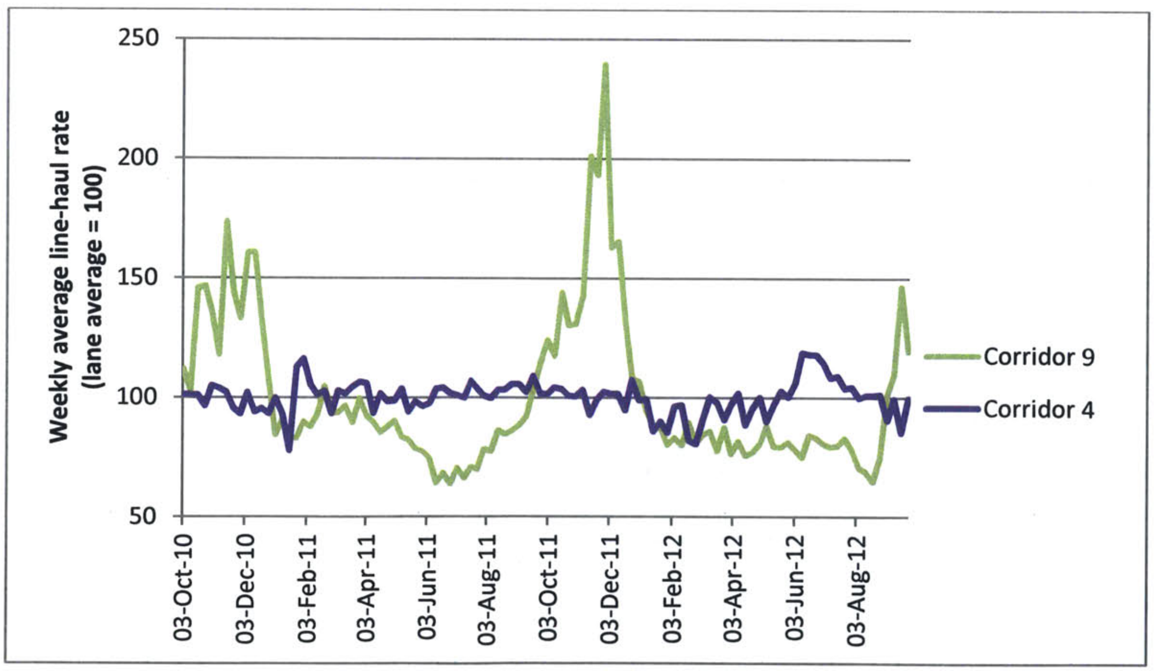

Figure 5.5: Comparison of range of average rates - corridors 4 and 9

5.3.3 Differences in seasonality of weekly average rates

Some of the corridors also exhibit seasonal patterns that differ markedly from one another. Figure 5.6 shows one particularly clear example of this phenomenon. Corridor 9 and Corridor 10 are a complementary pair of corridors. Corridor 9 has a repeated seasonal peak between May and August of each year. Corridor 10's peak occurs between May and August of each year.

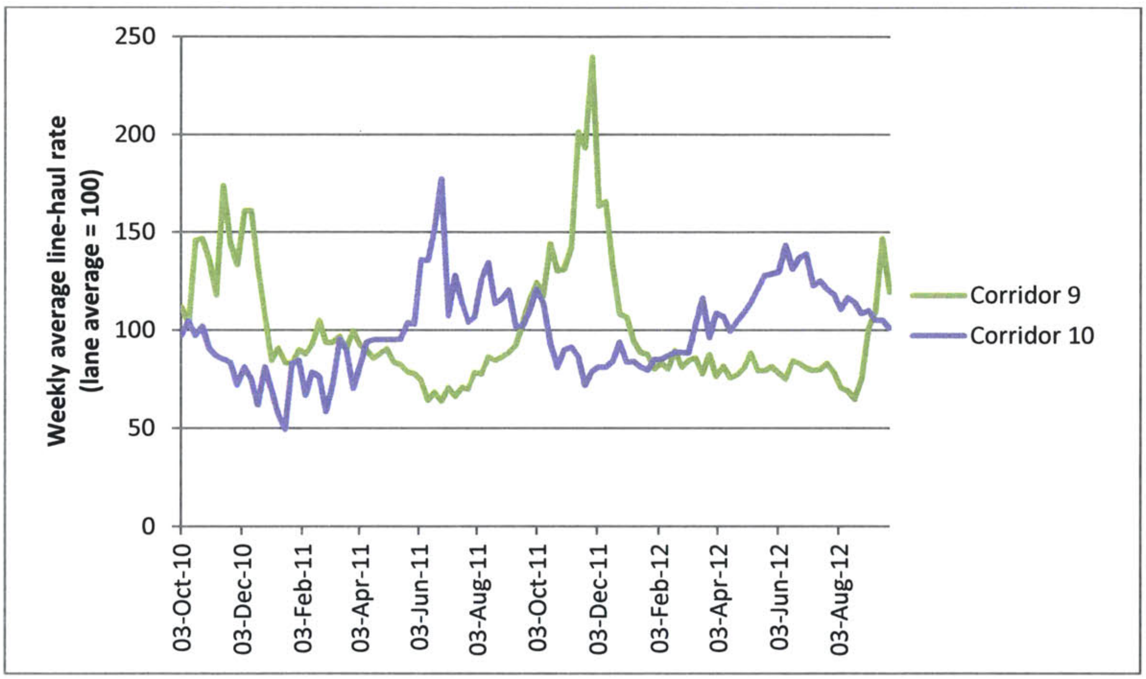

Figure 5.6: Timing differences of seasonal peaks

5.3.4 Differences in trends of weekly average rates

Finally, year-over-year rate trends were also examined. Average rates were calculated for both the first year and the second year of the study period and compared. Results are shown in Table 5.3. Changes in average rates ranged from a decline of 5.7% for Corridor 3 to a rise of 12.9% for Corridor 6. Even when examining only ten corridors, there are extremely significant differences in rate trends.

Table 5.3: Annual change in average rates

Corridor #	Region	Annual change in average rates
1	SE	-4.4%
2	SW	-2.4%
3	Midwest-SE	-5.7%
4	NE-Midwest	-2.7%
5	S	5.2%
6	NW	12.9%
7	Midwest	3.6%
8	E	-1.1%
9	W	4.4%
10	W	7.8%

5.4 Weighted index

The Baltic Indexes are calculated by summing the weighted rates of each of the component routes. A similar weighting approach could be used to develop a national level index for the LogisticsCo data. Like the Baltic Indexes, the ten corridors described in the previous sections were selected to represent a variety of geographies. Combined they represent only 3.9% of the entire volume in the dataset. Weightings are calculated based on the dollar value volume of each lane divided by the total volume of the ten corridors. Each weighting is rounded to be a multiple of 2.5% and shown in Table 5.4.

Table 5.4: Corridor weightings

Corridor #	Region	Weighting
1	SE	20%
2	SW	20%
3	Midwest-SE	2.5%
4	NE-Midwest	7.5%
5	S	10%
6	NW	10%
7	Midwest	7.5%
8	E	10%
9	W	5%
10	W	7.5%

The average weekly rates of each corridor are then weighted and summed to generate a weekly index. The result is shown in Figure 5.7. It shows a strong seasonal pattern with peak rates in June each year and the lowest rates occurring in January and February.

Figure 5.7: Weekly index generated from weighting ten corridors

The weekly index can also be simplified to a monthly index and compared with the LogisticsCo index developed from all data and the existing industry indexes. For clarity, only the DAT van rate index, the LogisticsCo all-data index and the LogisticsCo ten-corridor index are shown in Figure 5.8.

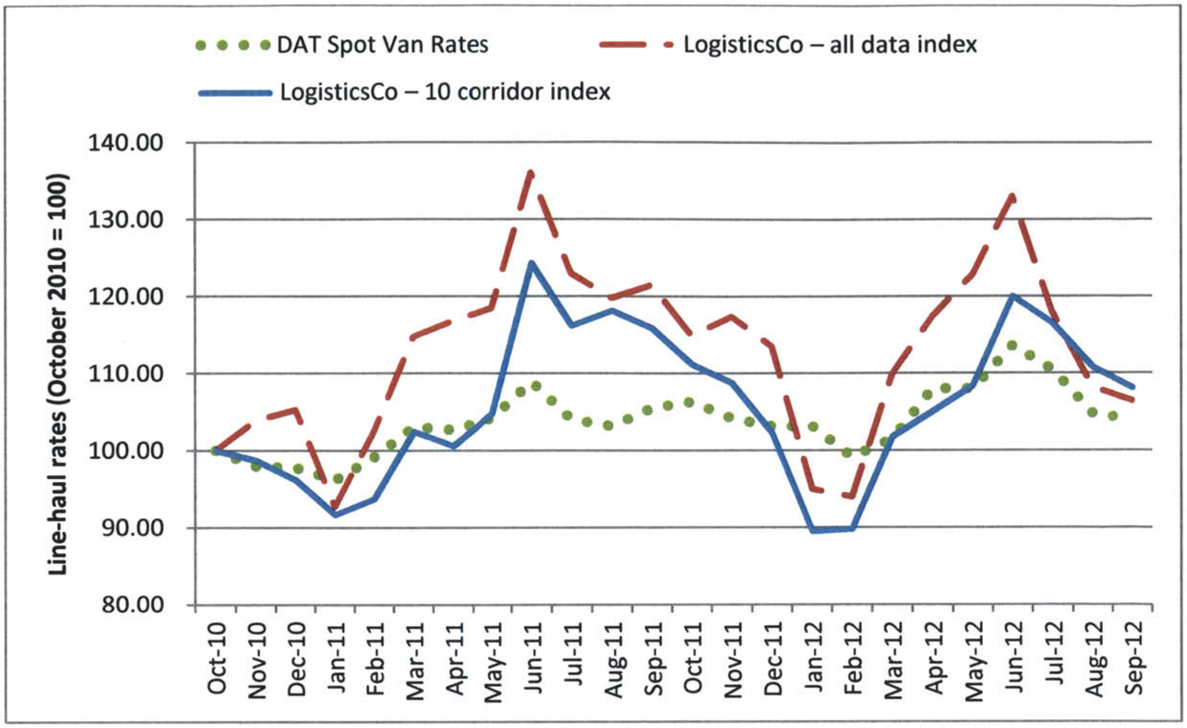

Figure 5.8: Comparison of LogisticsCo ten-corridor index with other indexes

With a correlation of 0.88, the index developed from the ten sample corridors tracks quite closely to the index developed from all LogisticsCo data. Although the ten-corridor index is generated from less than 4% of LogisticsCo's volume, it outperforms any of the industry indexes in representing the behaviour of the spot-market rates experienced by LogisticsCo. This suggests that a representative sampling and weighting method of corridors can perform well in the development of national level indexes.

5.5 Chapter summary

The purpose of the previous sections was to evaluate whether it would be necessary to calculate each individual corridor separately or whether the corridors were sufficiently similar that an aggregated national index could be effective. The ten sample corridors studied exhibited a wide variation in average rates and volatility. They also showed patterns of seasonality that differed dramatically. Finally, some corridors experienced rising rates over the period of the study, while others saw declining rates.

Any participant in the spot-market will be exposed to market conditions in a particular set of corridors. For a shipper with a simple distribution network, they may only be active in a handful of corridors. Larger, geographically diversified carriers and brokers may participate across nearly all corridors. However, even diversified market participants will have different concentrations of business and, therefore, will experience overall changes in rates in different ways.

This chapter provides evidence that the behavior of individual corridors can differ dramatically and suggests that indexes should be designed to provide rate information at the corridor level.

6. Temporal disaggregation – tender lead time and day of week

Prior research has identified tender lead time as having an influence on rates. That research also indicated that the day of the week that a load is picked up can also impact rates (Caldwell & Fisher, 2008). This chapter examines the role these factors play in the rates experienced by LogisticsCo. It then discusses the implications of these findings on the design of an index.

6.1 Lead time and day of week: OLS regression

Using ordinary least squares (OLS) regression, it is possible to test the relationship between lead time, day of week and rates. In OLS regression, the value of several independent variables can be used to calculate the value of the dependent variable. In this case, the dependent variable is the line-haul rate per mile.

Day of week was modeled with a series of dummy variables. Tuesday, Wednesday, Thursday, Friday, and Saturday were all established as variables. If a particular movement was loaded on a Tuesday, the variable "Tuesday" would have a value of 1 and all other day of the week variables would have a value of 0. To avoid perfect multi-collinearity, no variable for Monday was established. For a load that moves on a Monday, all of the day of the week variables would have values of 0.

Lead time was modeled in two different ways. Standard lead time is the number of calendar days between when the load is booked and when it is picked up at the shipper. Business lead time is the number of business days between when a load is booked and when it is picked up at the shipper. The objective of this differentiation was to determine which definition of lead time could better explain rate variability. In each case, lead time is defined as a continuous variable. Therefore, a lead time can be any positive number. For example, a lead time of 36 hours would be defined in the data as 1.5 days.

The performance of the regressions is evaluated using the R^2 term. In regression, this represents the proportion of the variability of the dependent variable that can be explained by the independent variables. It can range from zero to one.

Two corridors were selected to examine the effect of the variables in corridors with different characteristics. Corridor 1 was selected because it did not appear to demonstrate significant seasonality and was a relatively short-haul. Corridor 10 was selected because it did exhibit seasonality and was a long-haul.

6.1.1 OLS regression: Corridor 1

For reference the four-week moving average rate for Corridor 1 is shown below in Figure 6.1. There are no clear seasonal trends.

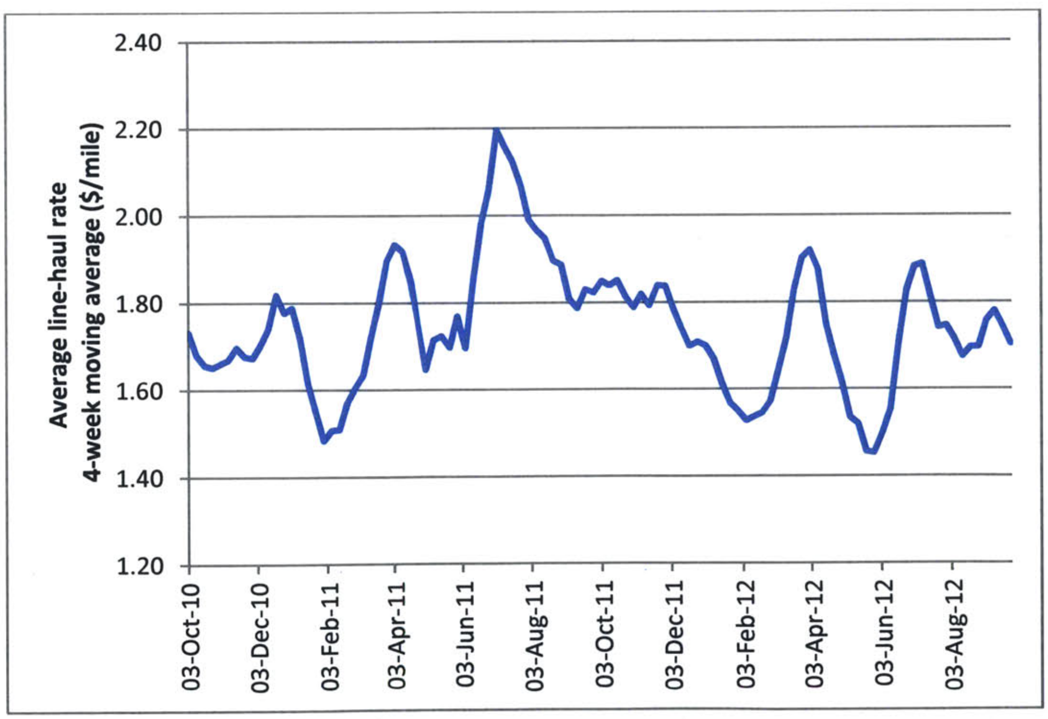

Figure 6.1: Average weekly line-haul rate – Corridor 1

The results of the model for Corridor 1 using standard, calendar day lead time are shown below in Table 6.1. Of note, Monday, Saturday and Sunday all had significant rate differences from the rest of the week. Despite the fact that all of the variables in this model are statistically significant, they are only able to explain a small proportion of overall variability in this corridor.

Table 6.1: Regression of standard lead time and day of week – Corridor 1

Variable	Coefficient	T-stat	p-val
Intercept	2.02	136.54	<.001
Lead time (days)	-0.19	-16.99	<.001
Lead time sq. (days^2)	0.014	9.41	<0.001
Monday pick-up	0.098	5.02	<0.001
Saturday pick-up	0.091	2.58	0.01
Sunday pick-up	0.25	5.96	<0.001
Fit Statistic			
R^2	0.178		

Table 6.2 shows the regression using business day lead time. It does explain slightly more of the overall variation in the dataset (R^2 of 0.192 vs. 0.178). Of interest, the Monday pick-up variable is no longer significant.

Table 6.2: Regression of lead time and day of week – Corridor 1

Variable	Coefficient	T-stat	p-val
Intercept	2.09	125.61	<0.001
Business Lead time (days)	-0.27	-17.90	<0.001
Business Lead time sq. (days^2)	0.026	9.57	<0.001
Saturday pick-up	0.11	3.20	0.001
Sunday pick-up	0.17	4.26	<0.001
Fit Statistic			
R^2	0.192		

6.1.2 OLS regression: Corridor 10

A similar analysis was completed for Corridor 10. For this corridor, there were no Sunday pick-ups. Figure 6.2 shows the four-week moving average rate for corridor 10 over the study period. This corridor exhibits seasonality and has peak rates between mid-May and late September in both years.

Figure 6.2: Average weekly line-haul rate – Corridor10

As shown in Table 6.3, the regression using standard lead times again showed a strong relationship between lead time and rates. This regression also suggested higher rates were associated with Monday and Saturday pick-ups.

Table 6.3: Regression of standard lead time and day of week – Corridor 10

Variable	Coefficient	T-stat	p-val
Intercept	1.78	57.01	<.001
Lead time (days)	-0.21	-8.25	<.001
Lead time sq. (days^2)	0.021	5.88	<.001
Monday pick-up	0.11	2.94	0.015
Saturday pick-up	0.21	2.03	.027
Fit Statistic			
R^2	0.119		

Table 6.4 shows the regression using business lead time. No days of the week were found to be significant when using this approach. The fit of this model slightly outperforms the fit of the model using standard lead times.

Table 6.4: Regression of business lead time and day of week – Corridor 10

Variable	Coefficient	T-stat	p-val
Intercept	1.86	53.68	<0.001
Business Lead time (days)	-0.29	-8.92	<0.001
Business Lead time sq. (days^2)	0.035	6.01	<0.001
Fit Statistic			
R^2	0.134		

For Corridor 10, there is strong seasonality and this provides an opportunity to further test the tender lead time relationship. Corridor 10 shows higher average prices in a summer season that runs between mid-May and late September each year. The data was separated into summer and non-summer periods and a regression was completed on each. This regression was in the same format as the previous analysis of the entire period that is summarized in Table 6.4. Results for each of the two periods are shown in Table 6.5.

Table 6.5: **Regression of business lead time by season – Corridor 10**

	Non-summer Season			Summer Season		
Variable	**Coefficient**	**T-stat**	**p-val**	**Coefficient**	**T-stat**	**p-val**
Intercept	1.53	45.20	<0.001	2.09	43.01	<0.001
Business Lead time (days)	-0.18	-5.91	<0.001	-0.26	-5.41	<0.001
(Bus. Lead time)^2 ($days^2$)	0.026	4.55	<0.001	0.025	3.00	0.003
Fit Statistic						
R^2	0.089			.145		

There are two key insights from this piece of the analysis. First, the relationship appears to hold in both periods of high rates and in periods of low rates. Second, the coefficient of Business Lead time is smaller in the segmented analysis than in the analysis that examines the entire two-year period. During the non-summer season, average lead time is 1.66 business days. This falls to only 1.32 days in the higher-rate summer season. Because the average rates are correlated with average lead times, a regression that includes data from the entire year will tend to overestimate the impact of lead time on rates. By analyzing each period separately, Table 5.8 provides a better understanding of the correlation of lead time and rates that is independent of seasonal factors.

6.1.3 OLS regression: All corridors

To confirm this relationship further, each of the 10 corridors was analyzed. In each case, a regression using business lead time days outperformed a regression using normal lead time days. The median increase in costs between next-day and same-day pick-up was 16.5% (high of 36%, low of 6%). The median difference between next day and two-day lead time was 6.5% (high of 13%, low of 1%). There were no further rate differences beyond three days.

Five of the corridors exhibited seasonality. For each of these, separate regressions similar to Table 5.8 were conducted. In four of five corridors, short lead times were associated with a greater increase in rates during peak season than during the non-peak season. The average increase for a same-day move over a next-day move was 19% during peak season and 11% during non-peak season. Figure 6.3 shows this relationship.

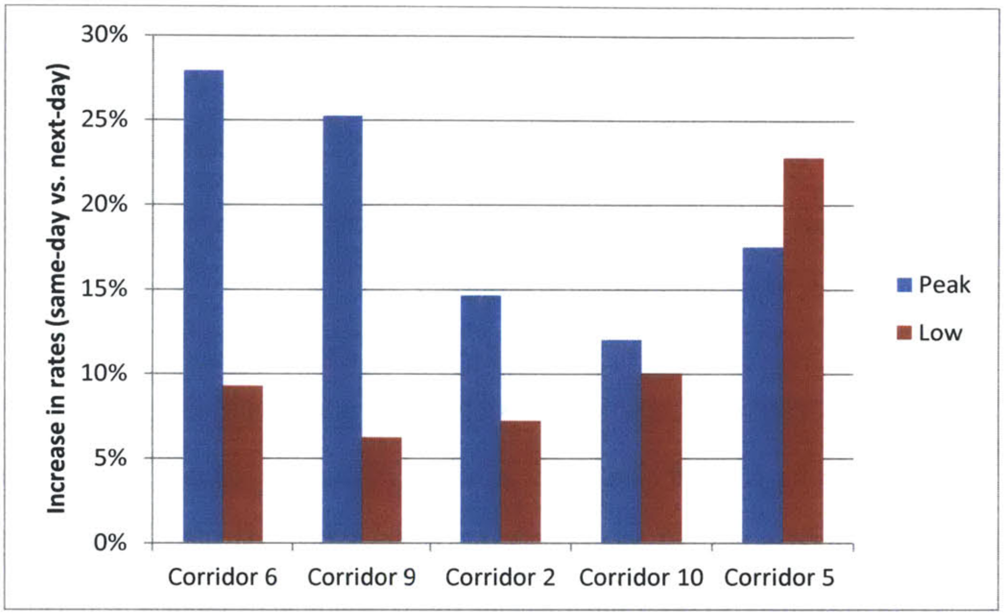

Figure 6.3: Increase in rates for a same-day move – corridors with seasonality

Corridor 5 does not behave in the same way as the other four corridors that experience seasonality. Of the five corridors, Corridor 5 has the smallest percentage difference between rates during the peak season and the slower season. In the peak season, Corridor 5's rates average $2.01 versus $1.69 during the slower season – a difference of 19%. By contrast, the other four corridors experience rate differences between peak and low seasons that range from 37% to 79%. It may be that the relationship between lead-time and seasonality requires more significant seasonal swings than those experienced in Corridor 5.

6.2 Temporal disaggregation: Implications

There is some evidence that day of the week of load pick-up is correlated with rates. In particular, Saturday and Sunday pick-ups exhibit statistically significant higher rates. This relationship has been noted in other studies (Abramson & Sawant, 2012, Caldwell & Fisher, 2008). Combined, Saturday and Sunday pick-ups represent less than 4% of all transported loads in the dataset. Given the relatively few loads that are affected, the most appropriate index design strategy is to make a specification that excludes weekend pick-ups.

It is clear that lead time is significantly correlated with rates in the studied lanes. Surprisingly, the number of business days of lead time provided a slightly stronger relationship with rates than did the total number of days of lead time. This may indicate that the relationship is driven more by business function behavior than the constraints of physically moving vehicles.

Lead time in this analysis could plausibly be acting as a proxy for other characteristics of the load that affect rates. Just as it takes longer for the brokerage to book loads during the summer season for Corridor 10, it is conceivable that it takes the brokerage longer to book loads that have difficult or challenging characteristics. These load characteristics might include special handling requirements, remote pick-up or drop-off locations or poor reputation of the shipper. While these characteristics do not appear as fields in a dataset, they would likely drive both higher rates and shorter lead times.

For the purposes of designing an index, the cause of the relationship is largely immaterial. The relationship exists and is quite substantial. The predicted same day rate (business lead time = 0.5 days) for Corridor 10 during the summer season is $1.97 per mile versus a next day rate (business lead time = 1.5 days) of $1.76 per mile.

By analyzing the relationship between lead times and rates in different portions of a corridor's seasonal rate cycle, it was found that rate premiums for short lead times was higher during peak season. This indicates that an industry participant with a high concentration of short lead time transactions will experience rate volatility higher than the industry average.

Similar to the geographical disaggregation discussed in Chapter 5, the relationship between tender lead time and rates suggests a second disaggregation is required. A market participant with greater exposure to transactions with short lead times will not experience changes in rates in the same way that a market participant with long lead times experiences rates. To reflect these differences, an index should specify what range of lead times the index represents. A logical choice in this situation would be to specify an index for next-day rates. 33% of all the loads in the dataset are booked with lead times of between 1 and 2 days. 38% of loads have lead times shorter than one day and 29% of loads have lead times of longer than 2 days. The next day rate represents a median lead time in the spot market.

6.3 Chapter summary

This chapter examined the impact of two factors related to time on rates in the spot market. It found that weekend pick-ups attracted premium rates, but that rates throughout the work week did not differ significantly from each other. It also found a significant and large relationship between tender lead time and rates. That relationship was also found to differ depending on the conditions in the market. During the higher-rate peak season short lead times were associated with much higher rate premiums than during the low-rate season. These results suggest that an index should specify a range of lead times for which the index rates are valid. The process of index specification necessarily limits the number of transactions that can be used to generate an index, potentially reducing its validity. The next chapter will examine this issue in more detail.

7. Inclusion of data from outside the index specification

One of the characteristics of indexes identified in Chapter 2 is that an index should be published frequently and that weekly publication is reasonable. Chapter 5 discussed how rates differed significantly based on the origin and destination of each load. Chapter 6 showed how rates differ significantly based on the amount of time between when a load is booked and when it is picked up by the carrier at the shipper's facility. It also suggested that specifying the index to represent transactions with next-day lead-times was appropriate. This process of index specification has consequences for availability of data as the number of transactions that occur on a specific corridor in any given week with a next day lead-time may be quite small.

A small sample size can lead to an inaccurate and volatile index. A larger data set with contributions from many industry participants may provide enough data for high-volume corridors. However, even with more data, many corridors will have low volumes. This chapter will examine two options for including information from transactions outside the specification of the index. Section 7.1 will examine the impact of including transactions with lead times of less than one day and more than two days. Section 7.2 will discuss the impact of including data from prior weekly periods.

7.1 Evaluating index performance

To evaluate the performance of different models of index generation, the set of transactions in a corridor will be divided into a training set and a testing set. Each transaction is given a random number between 0 and 1. All transactions with values less than 0.5 will be used to generate the index. The remainder of transactions will be used to test the effectiveness of the index. Corridor 5 was chosen as it exhibited some seasonal trends, some long-term trends and a wide distribution of rates that makes discerning underlying movements difficult. Figure 7.1 shows the distribution of rates of each transaction over the two-year study period.

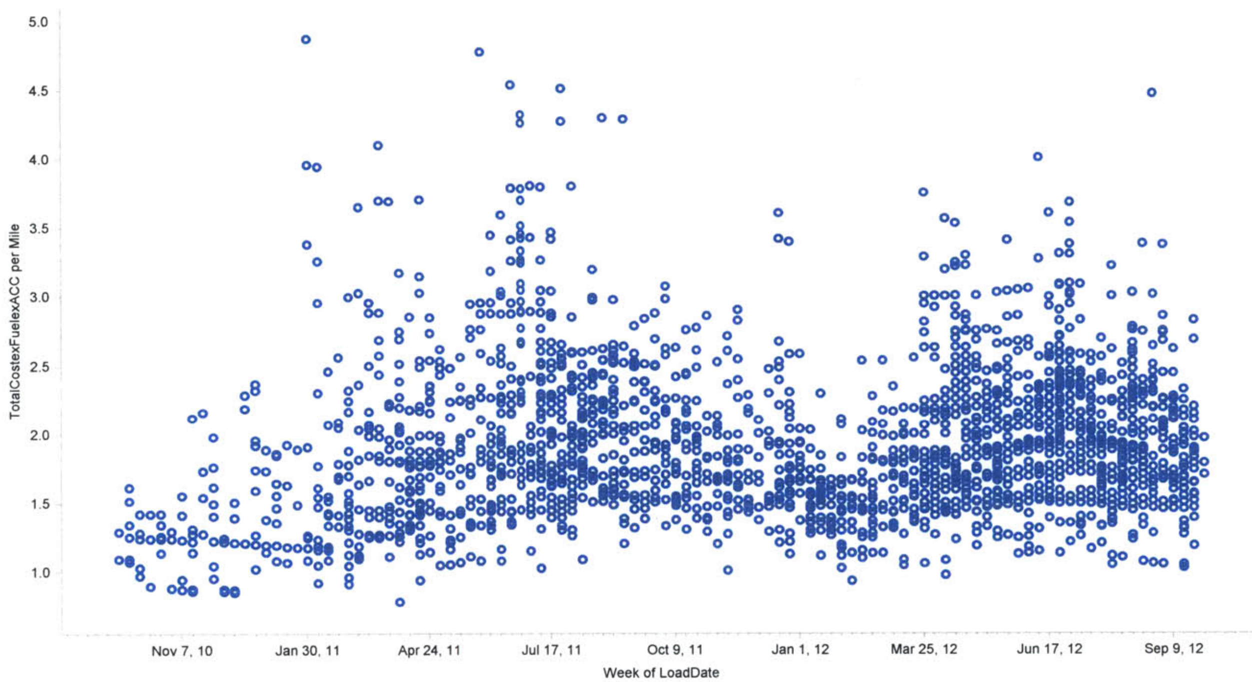

Figure 7.1: Corridor 5: Distribution of transactions by week and rate

As discussed in Chapter 6, there are dramatic differences between transactions with different lead times and an index can best handle this situation by specifying a range of lead times for which the index is valid. For testing purposes, the performance of the models will be evaluated solely on their ability to effectively predict transactions within the test set that have next-day lead times.

The index will calculate a weekly value. These values will then be compared to the values in the training set and errors calculated. Several statistical measures of fit will be used to compare the models. The mean percent error (MPE) is the average size of the error in terms of the observed values. It measures the bias of the model. If the value is 0, the model is unbiased. An observation other than 0, either positive or negative, indicates some bias. The mean squared error (MSE) and mean absolute percentage error (MAPE) both measure the accuracy of the predictions made by the index. MSE is calculated in the square of the units used in the index and does not normalize for the size of the observed values. MAPE measures the average absolute error in terms of a percentage of the observed values. Finally, the 'average absolute weekly change' is calculated. This is a measure of the volatility of the index and is calculated by taking the average of the absolute value of the weekly percentage change in the index over the two years. An index with high volatility is reacting to the inherent volatility in the sample data, but may not reflect changes to underlying trends. If accuracy is not impaired, an index with a lower absolute weekly change should be preferred as it is likely to be more useful by index users in their understanding of trends.

7.1.1 Model 1: Generate index from all transactions

The simplest model is to use all of the available data in the training set and calculate an average weekly rate. This method uses the maximum number of available values, but it is likely to be biased. This bias is generated by the relationship between lead time and rate. If the training set has a large number of same-day pick-ups, the average weekly rate will be biased upwards. Alternatively, if the training set has a large number of pick-ups with lead times longer than two days, the average weekly rate will be biased downwards. The training set has 1340 transactions.

7.1.2 Model 2: Generate index from only next-day transactions

To eliminate the tendency towards bias of Model 1, an index could also be generated by using only transactions with next-day lead times. This, however, significantly reduces the number of transactions used to calculate the index. In the training set, only 372 transactions had next-day lead times.

7.1.3 Model 3: Normalize the transactions based on lead time

Given the understanding of the relationship between lead times and rates discussed in Chapter 5, rates of all transactions can be normalized to a next-day equivalent. For example same day transactions are, on average, 16.5% more expensive than next-day transaction across the range of corridors analyzed. To calculate the next-day equivalent of a same-day rate transaction a multiplier is applied as shown below.

$$\text{“Next-day equivalent rate”} = \text{“Same-day rate”} * \frac{1}{1.165}$$

Likewise, transactions with lead times of two days or longer are, on average, less expensive by 6.5%. To calculate the next-day equivalent of a greater than two days lead time transaction, the multiplier below is applied.

$$\text{“Next-day equivalent rate”} = \text{“More than 2 days rate”} * \frac{1}{0.935}$$

More sophisticated normalizations are possible. An equation for lead time could be calculated and each transaction would have its own multiplier depending on the days, hours and minutes of lead time. However, for the purposes of this comparison, the simplified procedure is sufficient to demonstrate the impact on the index. Because each transaction is normalized, the full complement of 1340 transactions is used.

7.1.4 Model comparison

The three models were tested using the same dataset and the results are shown below in Table 7.1. As expected, the model that used all available data without normalization had a significant positive bias (11.3%). This bias was substantially reduced in both the index generated by “Only Next-day” data and the index generated on the data that had rates normalized by lead time. Another noticeable difference in performance between the indexes is the average weekly change in the index. This metric is important as it measures the volatility that a user of the index sees from week to week. The normalized data is less volatile than the model generated with only next-day data.

Table 7.1: Comparison of model performance

Model	MPE	MSE	MAPE	Average Absolute Weekly Change
All data	11.3%	0.176	31.8%	11.0%
Only Next-Day data	6.9%	0.190	32.3%	15.4%
Normalized All data	2.9%	0.155	28.9%	10.6%

The normalized data also performed best in terms of accuracy. MSE was reduced by 18.4% versus next-day data and 11.8% versus the index generated by all data. MAPE was also lowest for the normalized data. This result suggests that including the extra information provided by transactions with lead times outside the index specification is beneficial to the performance of the index when an appropriate normalization is used. Figure 7.2 shows a comparison between the index generated using "All Data" and the index generated using "Normalized Data" model. The bias of the "All Data" model can be clearly seen as it is consistently higher than the "Normalized Data" model.

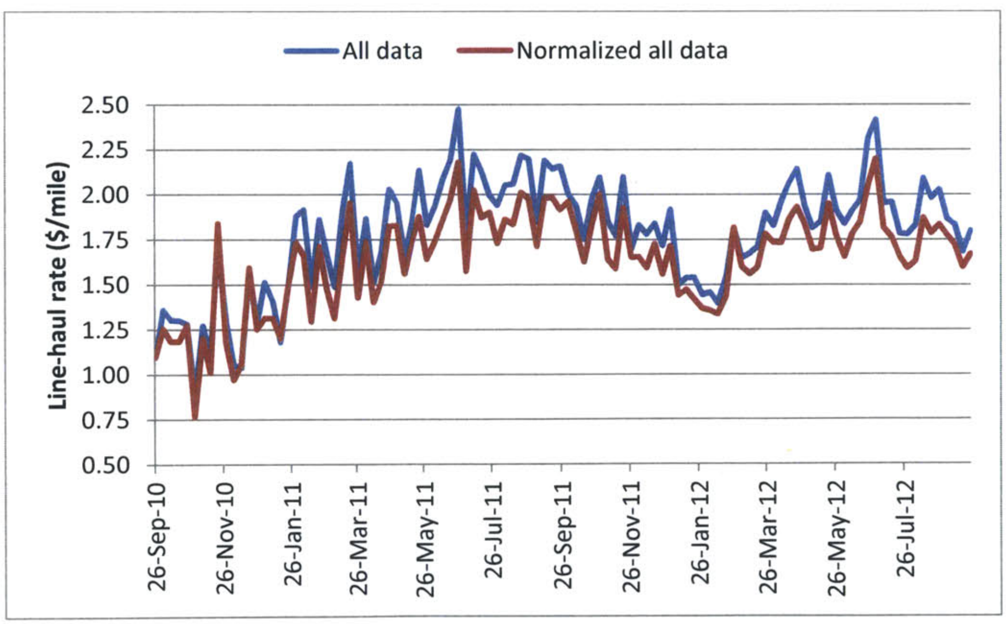

Figure 7.2: Comparison between All data and Normalized data models

7.2 Model smoothing

Each of the models developed in section 7.1 calculated the weekly index value solely on the observations that occurred during that week. These calculations ignore all of the transactions that have occurred in previous weeks. This approach can lead to significant volatility as each weekly value is produced from a relatively small sample of transactions.

An alternative method would be to calculate the weekly index using both past and current data. Two main approaches are common. An n-period moving average is calculated by summing each of the weekly index values and dividing by the number of periods. For example, if the index calculates values of 1.3, 1.6, 1.4 and 1.5 over a four-week period, than the moving average would be the sum of those values divided by four – 1.45. A moving average is most appropriate when long-term rates are level, but there is some statistical noise. Each value in the moving average period has equal weighting. In situations with trends, the moving average approach will have a pronounced lag.

A second method that can be used to incorporate data from previous weeks is exponential smoothing. In this method, a proportion of the calculated value of the index for the current week is based on new data, while the remainder is based on the value of the index from the previous week. This proportion is known as the smoothing constant. As an example, if I_t is the value of the index at time t, X_t is the average of the new data and a smoothing constant of α is used then:

$$I_t = \alpha (X_t) + (1-\alpha) I_{t-1}$$

Numerically, if the average of this week's new data was 1.5, last week's index was 2.0 and a smoothing constant of 0.7 is used than, this week's index is calculated as (0.7*1.5) + (0.3*2) = 1.65.

The moving average and exponential smoothing average were both tested using the normalized weekly data. The results are shown in Table 7.2. The standard model without smoothing is shown in bold for reference.

Table 7.2: Performance of different index smoothing options

	MPE	MSE	MAPE	Average Absolute Weekly Change
Normalized All	**2.9%**	**0.155**	**28.9%**	**10.6%**
4 week moving	3.8%	0.143	28.6%	3.0%
Exp smooth (α = 0.9)	3.0%	0.152	28.6%	9.1%
Exp smooth(α = 0.7)	3.2%	0.146	28.2%	6.7%
Exp smooth (α = 0.5)	3.4%	0.142	28.2%	4.6%
Exp smooth (α = 0.3)	3.2%	0.139	28.1%	2.7%
Exp smooth (α = 0.1)	0.1%	0.143	28.1%	1.0%

Overall, the results show that there are benefits to using a smoothing process in index calculation. Both the MAPE and MSE were reduced using a 4-week moving average and volatility is dramatically reduced. For exponential smoothing alternatives, accuracy was highest with a smoothing constant of 0.3. MSE was reduced by 10.3% versus the non-smoothed alternative. This indicates that there is value in including information from previous periods in the calculation of the weekly index. The index calculated in this way was better at predicting the values of the transactions in the test set. A normalized data model with smoothing constant of 0.3 performed significantly better than an index calculated only with next-day data. MSE was reduced by 26.8% and volatility as measured by weekly average change in the index was reduced by 82%.

While smoothing appears to improve overall accuracy and lower volatility, lower smoothing constants will increase the time it takes for an index to reflect major price changes in the market. Figure 7.3 compares the index with no smoothing with indexes with smoothing constants ranging from 0.1 to 0.5. The index with a smoothing constant of 0.1 clearly lags the other options in reflecting market trends.

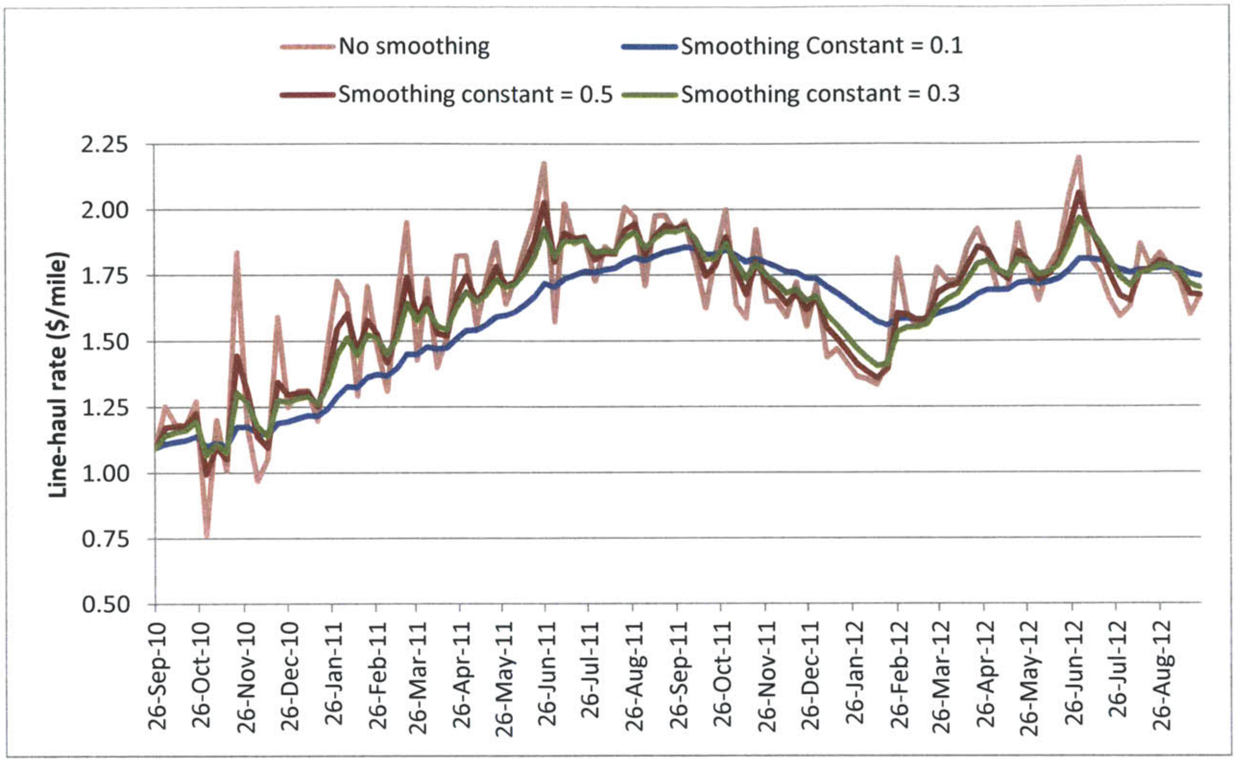

Figure 7.3: Lagging caused by smoothing process

The appropriate smoothing constant may change depending on the characteristics of the corridor measured. A corridor with more seasonal rate trends may be better served with an index with a high smoothing constant (or no smoothing at all) while an index of a corridor with more stability might be best represented using a low smoothing constant.

7.3 Overall model fit and implications for index uses

Sections 7.1 and 7.2 examined the accuracy of an index in representing the spot rates of next-day transactions on Corridor 5. It found that the index had a mean absolute percentage error (MAPE) of approximately 28% and varied slightly depending on the normalization and smoothing techniques used. This suggests that the rate of any individual transaction on a corridor is likely to differ significantly from the rates measured by the index.

The remainder of this chapter will examine the proportion of variability that can be explained by lead-time and underlying market changes. Understanding the extent to which variability in rates can be explained by these factors will then be used to discuss the usefulness of generating an index. It will refer back to the potential uses of indexes described in Chapter 2.

7.3.1 Analyzing overall fit

A linear regression model is used to calculate the proportion of variability that can be modeled by lead-time and underlying overall market conditions as represented by the index. The dependent variable is the line-haul rate per mile of each transaction in corridor 5 that is picked up on a weekday. The independent variables are lead-time (days), lead-time squared (days^2) and the value of the index ($/mile). Note that the index measures the average rate for a next-day move. Table 7.3 shows the results of this regression for Corridor 5.

Table 7.3: Analysis of overall index fit - Corridor 5

Variable	**Coefficient**	**T-stat**	**p-val**
Intercept	0.478	3.61	<.001
Business Lead time (days)	-0.443	-9.20	<.001
(Business Lead time)^2 (days^2)	0.079	7.05	<.001
Index ($/mile)	1.007	14.16	<.001
Fit Statistic			
R^2	0.226		

If the index were to be perfectly unbiased, its coefficient would be precisely 1. The result of this regression suggests that the index is relatively unbiased. The intercept coefficient can be interpreted to mean that a load with zero lead-time would pay $0.478 more per mile than the value of the index. As a comparison, a load with 0.5 days of lead-time would pay $0.28 more per mile than the rate indicated by the index. A load with a lead-time of 2.5 days would pay $0.14 less per mile than the index rate.

The important value to examine is the R^2 measure of fit. It suggests that only 22.6% of the rate variability in the data set can be explained by the factors in the regression. Combined with the result in section 6.2 that found the index had a MAPE of 28.8%, this suggests that an index does a relatively poor job of representing the rates of any individual transaction.

To confirm this finding, similar tests were completed on the other sample corridors and their measures of fit are shown in Table 7.4. Most corridor indexes were able to explain slightly more of the underlying variability than Corridor 5. Their accuracy, as measured by MAPE, was also slightly better. However, the majority of variability in all corridors remains unexplained.

Table 7.4: Analysis of overall index fit – all corridors

Corridor #	Region	R^2	MAPE
1	SE	0.268	17.3%
2	SW	0.385	15.1%
3	Midwest-SE	*Not enough data*	
4	NE-Midwest	0.046	25.1%
5	S	0.226	28.1%
6	NW	0.371	16.9%
7	Midwest	0.195	18.1%
8	E	0.252	12.8%
9	W	0.434	21.6%
10	W	0.402	15.5%
	MEAN	**0.286**	**19.0%**

7.3.2 Implications for index uses

Chapter 2 discussed four primary potential uses for indexes of the truckload spot-market: to provide price visibility and opportunistic optimization, to form the basis for derivatives, to use as an input in flexible-rate contracts, and to evaluate the costs of contracted arrangements versus participation in the spot-market. The inability of the index and other known factors to represent the variation in rates experienced in the market makes some of these uses unviable.

In particular, uses of the index that depend on the index representing a user's short-term experiences in the market will be unlikely to succeed. For example, one of the primary functions of derivatives is to allow a purchaser of a derivative to balance their exposure to changes in the real physical market against exposure to a financial derivative. According to the data in Table 7.2, the user's rates experienced in the real market will, on average, differ from the index by 19%. This limits the value of any hedging strategies involving index derivatives.

Using the index to inform short-term optimization strategies will also be extremely challenging. While the index measures the overall average rates in the market, the experience of any individual buyer or seller in the spot-market will differ significantly from the index. A carrier or shipper who shifts business from the contract market to the spot-market over the short-term based on the level of the index is likely to experience rates that differ significantly from their expectations.

As the index did represent an unbiased measure of the average rates over the study period, the index may still have value as information for decision-making and as a factor for inclusion in longer-term contracts. The index could still be used by shippers as they engage in procurement events. A shipper could use the index as an expected cost in the spot-market and compare that cost to contract rate bids received from carriers. The index rates could also be used as a negotiating tool when confirming contract rates.

Similarly, if a long-term contract covers a significant number of transactions, adding an index component to contracts could still be beneficial. It would increase contract rates when the spot market rates are higher overall and reduce contract rates when the sport market rates are lower overall. This would provide a mechanism for contract rates to change over time and reduce the incentives to hold frequent procurement events.

7.4 Chapter summary

This section examined the actual process involved in calculating an index value from a data set. The analysis shows that normalizing the rates of each transaction based on lead time generates a more accurate and useful index than alternative strategies. The analysis also indicates that a smoothing process that allows the index to incorporate information from previous periods led to a more accurate and stable index. This chapter also discussed the limited ability of an index to represent rates in the highly volatile spot market for truckload transportation.

8. Insights and conclusion

The purpose of this research was to examine a series of questions about rate indexes in the truckload spot-market. It reviewed the history of indexes in other similar industries and expanded upon a framework of characteristics that support successful freight indexes. It then examined existing trucking indexes within the framework and found that all lacked some of the necessary characteristics.

It also analyzed data provided by LogisticsCo, a non-asset based provider of logistics services, and found that they experienced much more seasonal variation than is expressed in any of the freight indexes. It found that rate behavior differed significantly between corridors, thus supporting index designs that disaggregate rates geographically. It also found that rates are strongly related to tender lead times and that the premium associated with short lead times differed by geography. Even within a single corridor, the premium associated with short lead times tended to be greater when market conditions were tight and rates were high. This supports index designs that disaggregate further based on lead times. No existing indexes provide that level of disaggregation.

This research also examined the process of determining an average weekly rate for a corridor with a specified range of lead times. It found that including transactions with lead times outside the specified range and using smoothing techniques improved the accuracy of the index and decreased its volatility. Finally, it reviewed the overall applicability of indexing for the spot-market. Indexes can only represent underlying rate trends and any individual transaction is likely to have a rate that differs significantly from the index. This rate volatility makes some proposed uses of a spot-market rate index improbable.

8.1 Management insights

Any individual market participant will likely find that their experiences with rates differ significantly from even a well-designed and valid index of spot-market conditions. However, indexes could be designed that provide rate information more accurately based on load characteristics. This paper notes specifically how indexes should be disaggregated by lead time as well as geography.

It also discussed how long-standing indexes like those of the Baltic Exchange were developed and continue to be supported by major participants in the industry. The DAT indexes share many of the characteristics of calculation and index design with those of the Baltic Exchange, but they are managed without significant industry participation. Major industry participants, particularly large brokerages, should determine if they want to pursue a spot-market index and evaluate whether DAT is the right partner.

8.2 Future research

This paper suggests strongly that accurate measurement of the spot-market truckload transport industry in the United States requires indexes that measure each corridor independently. However, with more than 10,000 total corridors and several hundred with significant volume, this data could be overwhelming. Future research that identifies groups of corridors that behave similarly based on their geographic and economic characteristics would be useful. This could lead to an understanding of benchmark or leading indicator corridors that could become key pieces of information for industry decision-makers.

Future research could also work to better understand the relationship between lead time and spot rates and how that relationship changes under different market conditions. The identified correlation is very significant and it's important to understand the causes of that relationship. Improved understanding in this area would have implications on decision-making for both shippers and brokers.

Reference List

Abramson, M., and Sawant, A. (2012). Impact of Risk Sharing on Competitive Bidding in Truckload Transportation. Thesis, Massachusetts Institute of Technology, Cambridge, MA.

Armstrong & Associates (2012). Top 30 Domestic Transportation Managers/Freight Brokers. Retrieved from https://www.3plogistics.com/Top_DTMs.htm

Caldwell, E., and Fisher, B. (2008). The Impact of Lead Time on Truckload Transportation Rates. Thesis, Massachusetts Institute of Technology, Cambridge, MA.

Caplice, C. (2007). Electronic Markets for Truckload Transportation. *Production & Operations Management*, 16(4), 423-436.

Caplice, C. & Sheffi, T. (2006). Combinatorial auctions for truckload transportation. In P. Cramton, Y. Shoham, & R. Steinberg (Eds.), Combinatorial auctions (Chapter 21). Cambridge and London: MIT Press

Gallu J., Brush S., and Fortado, L. (2012, June 28). Barclays Libor fine sends stocks lower as probes widen. *Bloomberg,* p. 1.

Grønvik, G. (2008). On commodity derivatives and the Norwegian initiatives to create a fish derivatives market. *Economic Bulletin – Norges Bank,* 79, 30-37.

Kavussanos, M.G. & Visvikis, I.D. (2004). Market Interactions in Returns and Volatilities between Spot and Forward Shipping Markets. *Journal of Banking and Finance,* 28(8), 2015-2049.

Kavussanos, M.G. & Visvikis, I.D. (2006). *Derivatives and Risk Management in Shipping, First ed.* Livingston, UK: Witherby Seamanship International

Kim, Yoo Joon (2013). Analysis of Truckload Prices and Rejection Rates. Thesis, Massachusetts Institute of Technology, Cambridge, MA.

Martens, B., and Suzuki, Y. (2012) Impact of Stale Rates. Research project. Iowa State University, Ames, IA.

Standard & Poors (2004). Transportation: Commercial. Industry reports.

Standard & Poors (2012). Transportation: Commercial. Industry reports.

Supply Chain Digest editorial staff (2010). Logistics News: Is New Bill Regulating Freight Brokers and Freight Forwarders Really Good for Independent Truckers? *Supply Chain Digest.* Retrieved from http://www.scdigest.com/ASSETS/ON_TARGET/10-06-22.php?cid=3542.

The Baltic Exchange. (2011). *Manual for Panellists: A Guide to Freight Reporting and Index Production.* London, UK.