Online pricing and capacity sourcing for third-party logistics providers

C. Lindsey[a], A. Frei[a], H. S. Mahmassani[a,*], H. Alibabai[a], Y.W. Park[b], D. Klabjan[b], M. Reed[c], G. Langheim[c], T. Keating[c]

[a]*Transportation Center, Northwestern University, 600 Foster Street, Evanston, IL 60208*

[b]*Department of Industrial Engineering and Management Science, Northwestern University, 2145 Sheridan Road, Evanston, IL 60208*

[c]*Echo Global Logistics, 600 W. Chicago Avenue, Chicago, IL 60654*

Abstract

Pricing shipments and sourcing capacity for a third-party logistics (3PL) provider operating in a spot market requires real-time decision making that is well suited for computer-based analytics. We outline a decision support system leveraging the 3PL's historical shipment data along with their knowledge of both sides of the shipping process to increase profits and to better perform the pricing and sourcing tasks. At the core of the system, are discrete choice models for shippers and carriers along with a profit maximization model. The discrete choice models predict the acceptance or rejection of an offer for a shipment to shippers and a bid for capacity to carriers. The profit maximization model determines the shipper price that maximizes the 3PL's expected profit. In addition to these models are procedures for determining a list of potential carriers for an incoming shipment and also for ranking those carriers. As its main outputs, the system produces a shipper price and a ranked carrier list. The system is applied to real-world data provided by a 3PL company with excellent results. The system is able to produce competitive yet profitable prices and to select potential carriers that would increase the 3PL provider's profits.

Keywords: Pricing, Sourcing, Optimization, Behavioral Modeling, third-party logistics, spot market

Introduction

Third party logistics (3PL) providers entered the U.S. logistics marketplace following deregulation in the early 1980s. On behalf of a company, a 3PL provider manages, controls, and delivers logistics activities ranging from the direct management of a company's supply chain to the sourcing of capacity for a single shipment in a spot market. The services offered by 3PL providers give firms the opportunity to lower their operating costs, gain greater flexibility in their supply chains, and to further concentrate on their core business activities. Additionally, 3PL providers are able to develop sophisticated systems and train personnel in the practice of logistics

*Corresponding author
Email: masmah@northwestern.edu, Phone: 847-491-7287

to a far greater extent than companies whose primary business is not logistics are willing or able to do, thereby lowering the client company's transportation costs. These, among other, potential benefits are what initially and continually make 3PL providers an attractive option to companies.

Different types of 3PL providers offer a broad range of services to companies, and are classified as asset-based or non-asset-based (Sheffi, 1990). Asset-based providers own assets (e.g., rolling stock, warehouses, etc.) that are first used to serve their primary clients and then sold as leftover capacity in the market place. Another business philosophy taken by some 3PL providers is to not own any traditional assets but to instead invest in systems and personnel as their primary assets. This leaves the company free to use other providers' assets and to leverage asymmetries in knowledge and expertise to enhance their profit. In this research, we take the perspective of the non-asset-based 3PL provider. Like all 3PL providers, these firms operate in an increasingly global, time-sensitive and dynamic environment in which market conditions can change rapidly. Because their systems, along with their expertise and knowledge of the market, are the primary assets by which non-asset-based 3PL providers survive, they need tools that help them navigate such a complex environment. This research develops the underlying methodology for a tool that addresses the need of a non asset-based 3PL provider to quote prices to a potential customer (i.e., a shipper) and source capacity (i.e., find a carrier) in spot markets.

A 3PL broker's knowledge and experience play an important role in the pricing and carrier selection decisions. An experienced broker has a good idea of average lane prices, and some knowledge of carriers likely to have capacity and willingness to serve a load within a given price range. However, the dynamic nature of the logistics market, gaps in experience between brokers, incomplete or poor access to information, and significant differences in market behavior across lanes can make quoting “good" prices and sourcing capacity challenging. Decision support systems guiding these tasks are valuable in this environment, providing a systematic approach for leveraging the large amounts of historical transaction information generally available in increasingly computer-supported environments.

We develop a comprehensive framework and decision support system that uses historical truckload shipment transaction data to price shipments and source capacity in spot markets with the ultimate objective of maximizing the 3PL provider's profitability. The objective is to recommend prices to quote a shipper in real time along with a list of potential carriers to contact to source the load. Though the presence of 3PL providers in the freight marketplace has been of considerable interest in recent years, the issues of pricing shipments and sourcing capacity have remained “below the radar” of academic research and the management science literature. Whereas most previous work regarding 3PL providers has centered on decisions of mode, shipment size, and carrier selection, this study contributes to the body of freight literature by examining a largely overlooked aspect of an important segment of the freight market operation; furthermore, this segment is growing in significance in light of continuing development and adoption of information and communication technologies.

This paper continues with a brief review of the 3PL provider literature. Section 2 then describes the structure of the modeling framework, and provides detailed derivations of the models that comprise the framework. In Section 3, the framework is applied to a case study using real-world data from a 3PL provider.

1. Literature Review

Much of the published literature addressing the 3PL provider sector has focused on assessing the extent to which 3PLs penetrate the freight market and their future prospects, the decision of firms to outsource logistics operations and the associated pros and cons, and also the factors affecting 3PL provider selection. Selviaridis and Spring (2007) provide a thorough review of the literature related to the role of 3PL providers in the business world. This vein of the literature explores the evolution of the relationship between 3PL providers and the firms they service, but it does not offer much insight into the operational practices of 3PL providers. The primary concern of our research is the pricing of shipments and sourcing of capacity by 3PL providers in spot markets.

Cheng and Qi (2011) investigated 3PL pricing using dynamic lot sizing models. In their study, Cheng and Qi determined the optimal pricing decision for a 3PL provider responsible for transporting raw materials from a supplier to a manufacturer where manufacturer demand is dynamic. They concluded that through greater cooperation, the 3PL provider could increase profits and the manufacturer could reduce transportation costs. This differs from our work in that our framework is intended as a real-time operational tool for spot markets, not as a planning tool. In addition, Cheng and Qi's work determined an optimal 3PL pricing policy within a single simple supply chain. We determine optimal pricing points for many different shipments that span many supply chains.

Regarding sourcing, most of the transportation procurement literature focuses on long-term contracts rather than spot markets. In the long-term contract literature, the sourcing problem is often conceptualized as an auction where shippers solicit bids from carriers for the service of a particular lane or group of lanes. Research on long-term contracts in transportation procurement have notably included topics such as auction design (Caplice, 1996), the historical use of auctions by shippers (Ledyard et al., 2002), carrier perspectives on auctions (Song and Regan, 2003), and the benefits of economies of scope and non-price variables in auctions (Caplice and Sheffi, 2003; Sheffi, 2004).

There is much less work regarding sourcing in spot markets. Figliozzi et al. (2003, 2007, 2008) modeled the spot market as a series of sequential auctions where a shipper or set of shippers offered shipments to a set of carriers, and carriers then bid on the shipments. The bidding problem is formulated as an equilibrium and decision theory problem. Their study conducted extensive simulation experiments to investigate the effect of the auction format, levels of sophistication for carriers, and learning capabilities of the carrier. The framework developed by Figliozzi (2004) is different from the 3PL business process presented in this research. In this work, the 3PL provider actively, rather than passively, seeks capacity, and is in fact the one effectively bidding for the shipper's load, albeit in a less formal process. However, in the limit, if the 3PL had complete information on the carriers' entire fleet and accepted loads, he would face a similar problem as the one addressed in Figliozzi et al.'s (2008) truckload pricing in a competitive environment.

Huang et al. (2011) consider the 3PL sourcing problem, in which a 3PL broker must actively seek capacity for a newly acquired shipment, and the key challenge is to determine which carriers to contact. They developed a network model of the relationships between 3PL brokers and carriers for a given set of loads. The network was then used to define metrics that correlate carrier selection strategies with 3PL profitability. Metrics included the total number of calls made and quotes given by a broker, among others. Simple linear regression was used to determine the correlation between the metrics and profitability. Huang et al. envision their model as a tool to inform operating policy guidelines and also to measure agent performance. Our modeling framework, however, operates in real time and considers attributes of the shipment and lane in addition to metrics of the 3PL-carrier relationship in the carrier selection process.

2. Model Overview and Derivation

2.1 Problem Description

The business process under consideration is the pricing of a shipment and the sourcing of capacity. Consider the following example: A shipper in need of a carrier to transport a shipment contacts a 3PL broker to facilitate this deal. The 3PL agent immediately (while still on the phone) gives a quote to the shipper, say $100. If the shipper accepts the quote, the 3PL provider takes responsibility of the shipment. Seeking to purchase capacity, the broker then proceeds to offer the shipment to several carriers. Some will agree to take the shipment for prices above $100, resulting in the 3PL provider potentially taking a loss; others will agree to take the shipment for less than $100, yielding a profit to the 3PL provider. The broker purchases capacity from the carrier that he or she expects to generate the most profit.

Pricing and sourcing are a freight broker's primary tasks and are critical to the success of a 3PL provider. Performing these tasks more intelligently represent real savings to the company in terms of time and money. The remainder of this section outlines the modeling framework and its mathematical models.

2.2 Modeling Framework

The modeling framework consists of three main component models: discrete choice models from both a shipper and carrier perspective, and a profit maximization engine from the perspective of the 3PL provider. Figure 1 summarizes the modeling framework and shows how each model interacts with the others. For a potential incoming shipment, the shipper choice model predicts whether a bid of a certain price (i.e. a quote) for that shipment from the 3PL to the shipper will be accepted. The shipper acceptance model uses attributes of the shipper, shipment, and lane among others to predict the decision outcome. Likewise, the carrier acceptance model predicts whether an offer of a certain amount for capacity for the newly acquired load from the 3PL to the carrier will be accepted. It uses attributes of the carrier, shipment, and lane among others to predict the most likely carrier decision. The profit maximization model uses the outputs of the two models as parameters of an objective function that determines an optimal price such that the 3PL provider's profit is maximized.

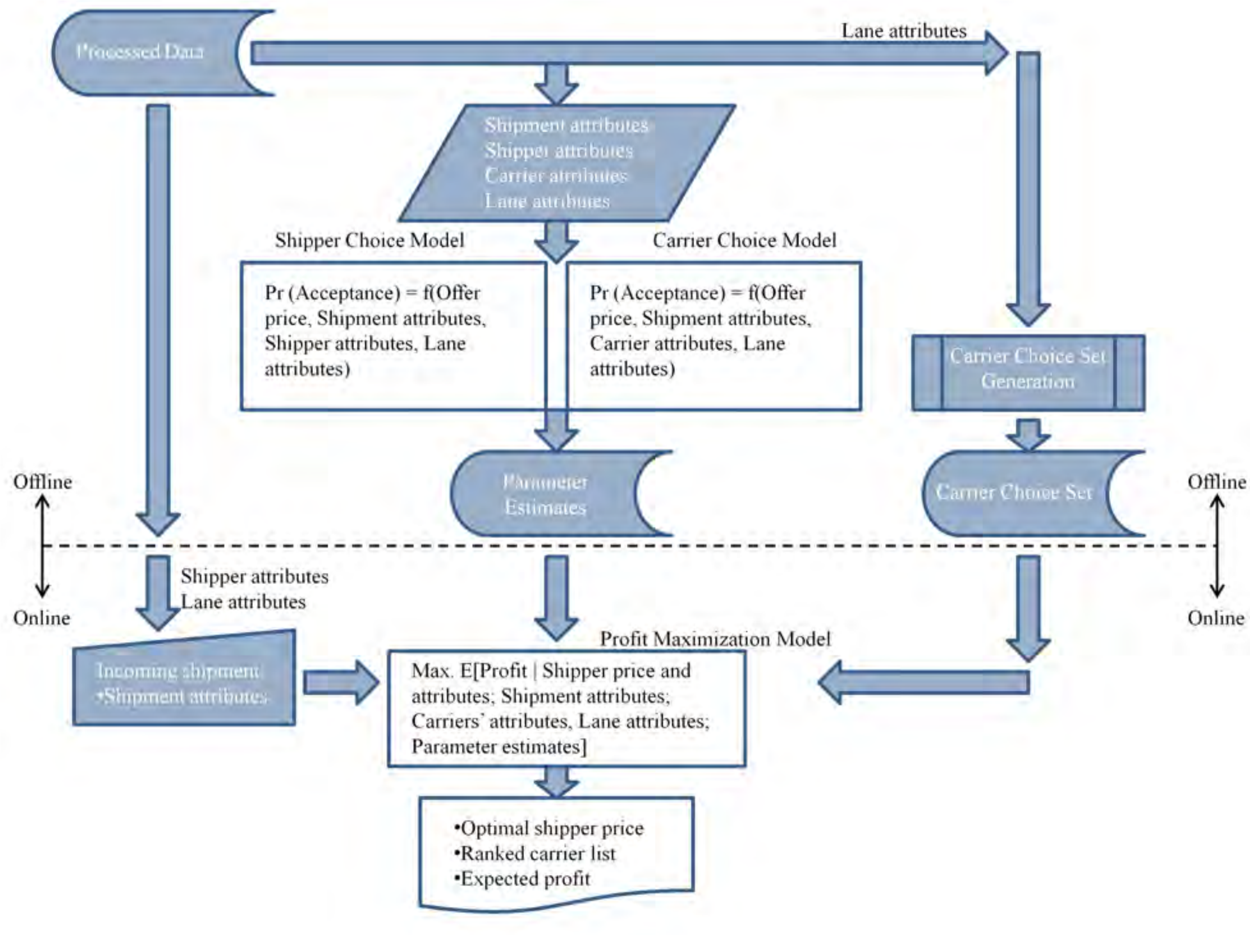

Figure 1. The modeling framework.

The core expected profit maximization engine of the framework drives the online operation of the system. It is triggered by the arrival of a request from a shipper to transport a load. Attributes of the shipment, along with the corresponding attributes of the shipper and lane (retrieved from the processed historical data), form the input to the model engine. As described mathematically in subsection 2.2.3, given the parameters of the other model components (the choice models) as well as user-specified parameters intended to control the degree of risk taking vs. risk aversion, the profit maximization engine determines an optimal price to quote to the shipper as well as a ranked list of potential carriers to contact for acquiring capacity.

The parameters of the choice models used in the real-time pricing engine are estimated adaptively, quasi-continuously reflecting experience acquired through the most recent experience. This is accomplished on a rolling horizon basis, triggered at either pre-specified intervals or according to a schedule dependent upon the 3PL provider's frequency and volume of transactions. As new data from recent transactions become available, older data is periodically removed from the stored processed data and replaced with more recent data.

The parameter estimation process is generally conducted off-line. It entails a certain amount of data filtering and processing, re-estimating the model coefficients, and forming carrier choice sets for each lane (i.e., the set of most promising carriers on each lane). These are relatively time-consuming activities that require analyst judgment, and are therefore appropriately performed offline.

The detailed formulation of the pricing engine and specifications of the respective choice models are described in the following subsections.

2.2.1 Behavioral Models

In this section we derive the discrete choice models used to estimate the acceptance of a bid (quote) for a shipment to shippers, and an offer for capacity to carriers. These models treat the acceptance/rejection decision as a function of price and other shipment, carrier and shipper attributes. The attribute coefficient estimates produced by these models are later used as parameters in the profit maximization model underlying the pricing engine.

Given a shipment i, i= *1,..,N*, let $Y_i = 1$ denote the acceptance by a shipper of the broker's quoted price, and $Y_i = 0$ rejection of that price (a similar binary variable is defined for the carrier decision). Consistently with random utility maximization, the we use the common logit form (McFadden, 1984) for the likelihood of acceptance, $Prob(Y_i = 1) = exp(X_i, \beta)/[1 + exp(X_i, \beta)] \equiv P(X_i, \beta)$ as the probability model, P. The vector β consists of the coefficients corresponding to each attribute in X_i (in a linear-in-parameters specification of the utility function), while b represents the estimates of β that maximize the likelihood function given below. Two choice probability models are specified in our framework, one for shippers and the other for carriers. In each model, the covariates, X, consist of the attributes that significantly affect shipper and carrier decision outcomes.

Because the data for this study, supplied by a 3PL provider, is "choice restricted," meaning that only one outcome of the binary choice variable, in this case price acceptance, was actually observed, we must take care to specify choice models that produce consistent parameter estimates. Choice restricted data are likely the norm among 3PL providers, and in many human-transaction driven environments. The value of recording rejections may not be readily apparent to 3PL agents, as deals are made in real time and data must be recorded manually after the deal is done. Hence changes to business as usual practices may be necessary in certain cases to obtain such data.

In the econometrics literature, so-called "choice-based" sampling is a common phenomenon and is normally treated with a correction factor for the choice-based bias. The extreme case, in which the observed data is "choice-restricted" is not as often discussed. Cosslett (1981), and later Steinberg and Cardell (1992), suggest using a supplementary sample to augment the likelihood function and render the model identifiable. A supplementary sample is an unbiased sample containing information on the independent variables in the model (e.g., origin, destination, lead time, etc.), but no information on the dependent variable, in this case acceptance/rejection. The supplementary sample together with a known share of the sub-population that accepts, $Y = 1$, facilitates the derivation of a decomposed log-likelihood function that can be used to consistently estimate binary choice models.

For a random sample of size N with data on the binary response variable, Y, and covariates, X, the conditional log-likelihood function can be written, following Steinberg and Cardell (1992), as follows:

(1) $\mathcal{L}(b) = \prod_{i=1}^{N} P_i^{Y_i}[1 - P_i]^{1-Y_i}$

(2) $\mathcal{L}\mathcal{L}(b) = \log \mathcal{L}(b) = \sum_{i}^{N} Y_i \log P_i + (1 - Y_i)\log(1 - P_i)$.

In order for the model to be identifiable, Y must have been observed and $Var(Y) \neq 0$. We use the supplementary sample to introduce variation in Y. To do this, the log-likelihood function (2) is decomposed into two parts, with the assumption that the supplementary sample represents the whole population and the choice-restricted sample contains all cases with $Y = 1$:

(3) $\mathcal{L}\mathcal{L}(b) = \sum_{i=1}^{N} \log(1 - P_i) - \sum_{\substack{i=1 \\ i:Y_i=1}}^{N} \log P_i - \sum_{\substack{i=1 \\ i:Y_i=1}}^{N} \log(1 - P_i) = \sum_{i=1}^{N} \log(1 - P_i) +$

$\sum_{\substack{i=1 \\ i:Y_i=1}}^{N} \log(P_i/(1 - P_i))$.

A related pseudo-likelihood can be applied to a pooled sample with the following notation: r_0= the sampling rate in the supplementary sample; r_1 = the sampling rate in the choice-restricted sample; $N = r_0 T$, the size of the supplementary sample, where T is the size of the sub-population (drawn from the infinite population); $M = r_1 sT$, the size of the choice-restricted sample, where s is the share of the sub-population with $Y = 1$; and $I = N + M$, the total sample size.

For pooled samples with general r_0 and r_1 we replace (3) with:

(4) $\mathcal{LL}^*(b) = \sum_{i=1}^{N} \log(1 - P_i) + \frac{r_0}{r_1} \sum_{i=N+1}^{N} \log(P_i/(1 - P_i)).$

Differentiating (4), the gradient and the Hessian for the pseudo-log-likelihood $\mathcal{LL}^*(b)$ are:

(5) $g(b) = -\sum_{i=1}^{N} P_i X' + \frac{r_0}{r_1} \sum_{i=N+1}^{I} X'_i$

and

(6) $H(b) = -\sum_{i=1}^{N} P_i (1 - P_i) X'_i X_i,$

respectively. Equation (4) can be maximized using a number of algorithms. In this study we use the BFGS algorithm (Broyden, 1970; Fletcher, 1970; Goldfarb, 1970; Shanno, 1970).

It is important to note that though we derive and implement the logit model in our system, the flexibility of the framework accommodates any choice model form (e.g., probit, mixed logit, etc.). Also, the log-likelihood function specification employed is solely driven by the choice-restricted nature of the data. Because choice-restricted data is likely the norm for much of the logistics industry, the log-likelihood specification used here may be necessary to employ our system throughout the 3PL provider community. In the event that a 3PL provider does have data on both accepted and rejected bids and offers, more common log-likelihood functions would be used.

2.2.2 Carrier Choice Set Definition Procedure

The objective of this step is to identify the best candidate carriers for a shipment and then rank them. The procedure assumes that the most frequent and the most currently active carriers on that particular lane are better candidates for accepting a load at a reasonable price. Those carriers are identified and placed into the choice set. Each potential carrier is then ranked according to the expected profit they are likely to yield to the 3PL for the already accepted shipper quote, using the carrier acceptance model described in the preceding section. In lanes where there are not many active carriers, the choice set is enriched with carriers from nearby lanes. We consider a lane to be nearby if its origin and destination is geographically close to the origin and destination of the lane in question. The choice set generation procedure can also be adapted to incorporate additional information on carrier availability in or preference for particular lanes.

2.2.3 Profit Maximization Model

The estimated coefficients (b) for the different covariates (X) are used to formulate the profit maximization problem as a function of price alone, as the remaining covariates are set to the known attribute values of the shipment under consideration. In other words, the estimated coefficients of the choice models that operate on (multiply) price are collected (i.e., summed) into one term, while all other components (consisting of the known attribute values multiplied by their respective coefficients in b) are collected into another term – effectively a slope and intercept. For shippers, the collected term for price-related parameters (i.e., slope) is denoted by m, whereas the collected term for non price-related components (i.e., the intercept, since those terms are already known in the profit maximization formulation) is denoted by q. Likewise, for carriers, the collected term for the price-related parameters is denoted by o_k, whereas the collected term for the non-price related components is denoted by v_k, where the index k denotes the carrier.

Note that the carrier slope and intercept terms have an index while the shipper slope and intercept terms do not. This is because, at the point at which the profit maximization routine is invoked the shipper and all its attributes are known, while the carrier that will actually transport the shipment is unknown, as there are usually numerous potential carriers. This is reflected in the optimization model formulated below. The following notation is used throughout the remainder of the paper:

k: carrier k, $k = 1, \dots, K$;

m: collected term for a given shipper's price-related parameters (i.e., shipper slope);

q: collected term for a given shipper’s non-price related components, each consisting of the product of covariate value and corresponding coefficient, (i.e., shipper intercept);

o_k: collected term for potential carriers’ k price-related parameters (i.e., carrier slope);

v_k: collected term for potential carriers’ k non-price related components (i.e., carrier intercept);

$S_{Acceptance}$: overall rate at which shippers accept offers from the 3PL provider (estimated from a sample of data taken for this purpose);

$\Psi(.)$: acceptance probability function that predicts the given shipper's acceptance of an offer from the 3PL provider;

$\Gamma_k(.)$: acceptance probability function that predicts a potential carrier k's acceptance of a bid for capacity from the 3PL provider.

The profit optimization model determines the quoted shipper price, ψ (the decision variable), that maximizes the 3PL provider's expected profit. The profit for the 3PL provider is the difference between the offer accepted by the shipper and the bid accepted by the carrier. However, at the time of the quote neither the offer nor the bid price which will be accepted by the best carrier are known. In addition, we also do not know which carrier will present the best opportunity for the 3PL provider. We make the assumption that a carrier in the carrier choice set accepts the minimum bid, i.e., carriers accept at the lowest price at which they are willing to service the load.

We are given the shipper's acceptance probability function $\Psi(\psi) = exp(m\psi + q)/1 + exp(m\psi + q)$ and carrier k's acceptance probability function $\Gamma_k(\gamma_k) = exp(o_k\gamma_k + v_k)/1 + exp(o_k\gamma_k + v_k)$ for all carriers $k = 1, \dots, K$, where ψ is the shipper price and γ_k is the carrier price (see Figure 2 (a)). The probability density function of the price accepted by a given carrier is the first derivative of the logit function and is expressed as:

(7) $$f_k(\gamma_k) = \frac{d}{d\gamma_k}\frac{exp(o_k\gamma_k+v_k)}{1+exp(o_k\gamma_k+v_k)} = \frac{o_k \cdot exp(o_k\gamma_k+v_k)}{(exp(o_k\gamma_k+v_k)+1)^2}.$$

The expected profit given ψ should consider all possible bids for capacity to carriers. We assume that the lowest bid is accepted. Hence, the objective function is defined as

(8) $E[\text{Profit}|\psi] = \Psi(\psi) \int \dots \int_{-\infty}^{\infty} [\psi - \min(\gamma_1, \dots, \gamma_K)] \prod_{k=1}^{K} f_k(\gamma_k) d\gamma_1 \dots d\gamma_K$

and the optimization problem is formally written as

(9) $\max_{\psi \geq 0} \mathrm{E}[\text{Profit}|\psi]$.

The optimization problem is visually described in Figure 2.

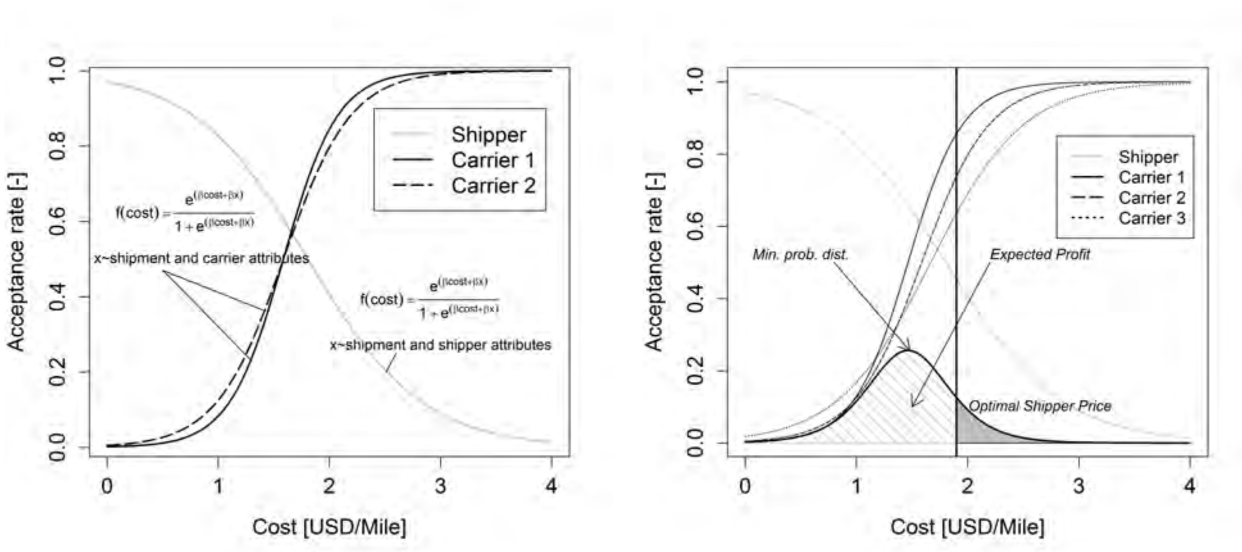

(a) Acceptance probabilities for a given shipment (b) Expected profit and the resulting optimal price

Figure 2. Optimization inputs and outputs.

This optimization problem is challenging to solve within a few seconds of running time, a necessity since the system operates in real time. The integral in Eq. 8 does not admit a closed form solution; numerical integration methods were not found to be satisfactory, especially compared to Monte Carlo sampling, which proceeds as follows:

1. Using the inverse CDF method, we generate a single sample of size k (one observation for each carrier) of carrier prices. In our case the inverse of Ψ yields $\mathrm{p_k} = -\log\left(-\frac{u_k \cdot exp(o_k) - exp(o_k)}{u_k}\right)/o_k$, where u_k is a uniform $[0,1]$ sample.
2. From this sample, we select the minimum price: $\mathrm{p_{min}} = \min\{\mathrm{p_1,p_2,\dots,p_k}\}$.

3. We repeat this process t times to produce a distribution of minimum carrier prices: $(\mathrm{p_{min}})_{\mathrm{j}}$ where $j = 1, \ldots, t$.

The optimization problem over the samples becomes $\max_{\psi} \frac{1}{t} \sum_{j=1}^{t} F(\psi)\left[\psi - (\mathrm{p_{min}})_{\mathrm{j}}\right] = \max_{\psi} \frac{t \cdot \psi - \sum_{j=1}^{t} (\mathrm{p_{min}})_{\mathrm{j}}}{t\left(1+exp(-\{m\psi+q\})\right)}$. This model maximizes expected profit while not considering risk. This could lead to a situation where unreasonably high prices are quoted to a shipper because there is a small chance of achieving a very significant profit. Such a myopic approach could lead to an erosion of market share, loss of customer goodwill, and have dire long-term consequences. For these reasons we incorporate risk into the model.

Let $w(\psi)$ be a weight that depends on ψ and allows the 3PL provider to tradeoff the profit and the risk of losing deals. We update the model by adding this weight to the objective function, which is now specified as:

$$\max_{\psi} \frac{1}{t} \sum_{j=1}^{t} w(\psi) F(\psi)\left[\psi - (\mathrm{p_{min}})_{\mathrm{j}}\right] = \max_{\psi} \frac{w(\psi) \cdot t \cdot \psi - \sum_{j=1}^{t} (\mathrm{p_{min}})_{\mathrm{j}}}{t\left(1+exp(-\{m\psi+q\})\right)}.$$

If $w(\psi) = 1$ for every ψ, then the model strictly maximizes profit. It can be calibrated so that the overall acceptance rate (and therefore the risk) is generally higher, lower, or equal to the historical acceptance rate. We propose the weight function $w(\psi) = \left(F(\psi) - S_{Acceptance} + 1\right)^{exp(\alpha)}$, where α is a parameter to be calibrated. This optimization problem can be solved numerically by setting the first derivative to zero and employing a line search algorithm:

(10) $$\frac{d}{d\psi} \frac{w(\psi) \cdot t \cdot \psi - \sum_{j=1}^{t} (\mathrm{p_{min}})_{\mathrm{j}}}{t\left(1+exp(-\{m\psi+q\})\right)} =$$

$$\frac{\left(1 - S_{Acceptance} + \frac{1}{1+exp(-\{m\psi+q\})}\right)^{exp(\alpha)}}{1+exp(-\{m\psi+q\})} +$$

$$\frac{m\left(t \cdot \psi - \sum_{j=1}^{t} (\mathrm{p_{min}})_{\mathrm{j}}\right) \cdot exp(-\{m\psi+q\}) \cdot \left(1 - S_{Acceptance} + \frac{1}{1+exp(-\{m\psi+q\})}\right)^{exp(\alpha)}}{t\left(1+exp(-\{m\psi+q\})\right)^{2}} +$$

$$\frac{m\left(t \cdot \psi - \sum_{j=1}^{t} (\mathrm{p_{min}})_{\mathrm{j}}\right) \cdot exp(-\{m\psi+q\}+\alpha) \cdot \left(1 - S_{Acceptance} + \frac{1}{1+exp(-\{m\psi+q\})}\right)^{exp(\alpha-1)}}{t\left(1+exp(-\{m\psi+q\})\right)^{3}}.$$

2.2.4 Carrier Ranking Procedure

Recall that after a price has been quoted to and accepted by the shipper, a score to each carrier is provided so that calls to carriers may be prioritized. Carriers are ranked according to the likelihood that they will accept a price that generates a profit for the 3PL. Given the shipper's accepted price, ψ^*, derived from (9), an upper bound on the 3PL provider's profit range is yielded. The shaded area in Figure 2 (b) represents the price range for which the 3PL provider is profitable. Hence, their chance of making profit from carrier k can be expressed as

(11) $$P_k[\text{price} < \psi^*] = \int_0^{\psi^*} f_k(\gamma_k) d\gamma_k.$$

We directly use probabilities $P_k[\text{price} < \psi^*]$ for $k = 1, \dots, K$ to rank carriers.

For a given shipment, the carriers ranked in this procedure are those previously generated by the carrier choice set procedure. By contacting carriers in this order, we balance a broker's need to quickly find capacity and also to generate a profit for the company.

3. Case Study

The data for the study are truckload shipments managed by a U.S.-based 3PL provider operating in the U.S., Canada and Mexico. The data spans the years 2005 to the end of 2012 and contains approximately 556,000 shipments. Each observation contains the following: time of the shipment, origin and destination of the load, equipment type, number of stops in the tour, the price accepted by the shipper, the cost charged by the carrier, and a unique shipper and carrier identification. In addition to shipment data, data on carrier preferences (e.g., a preferred service region) and carrier availability is used to enrich the shipment data and aid in the specification of the choice-models.

3.1 Data Cleansing and Delineation

The objective of the data delineation and cleansing process was to generate a set of data for analysis that allows a clear interpretation of the effects of various attributes. In accomplishing this, it was necessary to maintain the natural variation in the data while ensuring that all of the observations are consistent with one another. This was achieved by filtering the raw data on several criteria including the geographic boundaries of the shipments, mode, and distance among others.

First, we delineated and retained a year's worth of the most recent shipments with an origin and destination inside the contiguous U.S. The international shipping process is inherently different than the domestic process and would require an entirely different modeling framework. After that, the data was delineated to exclude short-distance trips. Like international shipping, short-distance trips represent a different market, such as drayage, and would require an alternate model perspective. Next, only records for basic shipping services with complete cost information were kept. 3PL providers sometimes provide additional services beyond basic shipping, such as loading/unloading, the nature and extent of which could greatly affect pricing decisions and shipper/carrier behavior. Lastly, outliers were removed based on price-per-mile values that were considered unreasonable.

3.2 Data Mining and Variable Creation

Once cleansing and delineation were complete, the data was spatially clustered so that geographic effects in the brokerage process are more evident. Spatial attributes of a load are important determinants of its price and acceptability to carriers. Certain geographic areas command lower or higher prices depending on the level of activity within and between those areas. This information can help to build an understanding of the spatial distribution of shipments, variations in shipment prices/costs and frequency among lanes, and serve as a basis for defining many variables to be used in the choice models. Since there are not enough observations at the geographic level of traditional lanes (i.e., metropolitan area to metropolitan area), we cluster the data. Clustering was done via the K-means algorithm (Hartigan et al., 1979). The K-means algorithm was chosen because it easily captures both the spatial and temporal (i.e., frequency of shipments) aspects of the data in the clustering process. With K-means, each point is weighted proportionally to its frequency as an origin or destination for a shipment in developing clusters.

The clustered data was then used to create variables for the behavioral models. There are four primary variable types included in the analysis: shipment, shipper, carrier and lane variables. Each variable type identifies the attributes that most affect acceptance or rejection of the offer/bid. For instance, shipment attributes include information such as lead time, number of stops, and distance, among others. Shipper and carrier attributes include how often they use the

3PL provider's services and what range of prices they normally accept. Lane attributes include estimates of the volume of traffic and the number of carriers that operate on the lane.

Variables are defined on a rolling horizon basis. In other words, the measurement of each variable is taken at the time, or some period up to that time (e.g., the previous 30 days), of the actual transaction. Thus, in calibrating the choice models, the coefficients are a true reflection of the influence of that variable at the time the shipment was made. For further information on the data mining and variable creation process, the reader is referred to Lindsey et al. (2013).

3.3 Results and Discussion

Next, we present and discuss the results of the discrete choice and profit maximization models. In general, we discuss our results in reference to what is expected in shipper/carrier behavior for the choice models, and the observed historical prices regarding the profit maximization model.

3.3.1 Choice Model Results and Discussion

With the choice models, the goal is to accurately measure the relative acceptance probability of a given shipment as a function of the predictors including price, shipment attributes and shipper/carrier attributes. In the calibration process, several model specifications with varying amounts of clusters and different rolling horizons for the data (i.e., 3 months of data, 6 months of data, 9 months of data, etc.) were tried. The final cluster count and time horizon are the ones which provide the best predictive capability and that fit the operational constraints of the 3PL provider. As is common practice, the variables that were the most statistically significant decided the final specification.

In specifying behavioral models for the shippers and carriers, we chose variables that captured shipment, shipper, carrier, and lane attributes. The model is largely defined by how these variable types affect price, which is the primary decision driving acceptance or rejection. Also, price is the only variable over which the 3PL provider has control. Because of this, variables that include price change the slope of the behavioral curves while variables that do not include price only affect the intercept. Figure 3 shows an abbreviated version of the calibration results. Full calibration results are included in the appendix.

An important variable in the behavioral models that warrants further discussion is the “MinDist" variable. Often, carriers do not price shipments according to per-mile rate schedules for relatively short distance trips. Instead, they use a flat price. This practice may cause price-per-mile measures for short distance trips to exhibit a large degree of variation and limit the statistical power of our models. We control for this with a minimum distance indicator variable that identifies short distance trips (defined as those that are less than 300 miles). This threshold was chosen upon conducting an analysis for determining breakpoints in data using regression models. The procedure determined that a breakpoint occurs at approximately 292.4 miles. Interested readers are encouraged to refer to Lindsey et al. (2013) for further discussion of the results of this analysis.

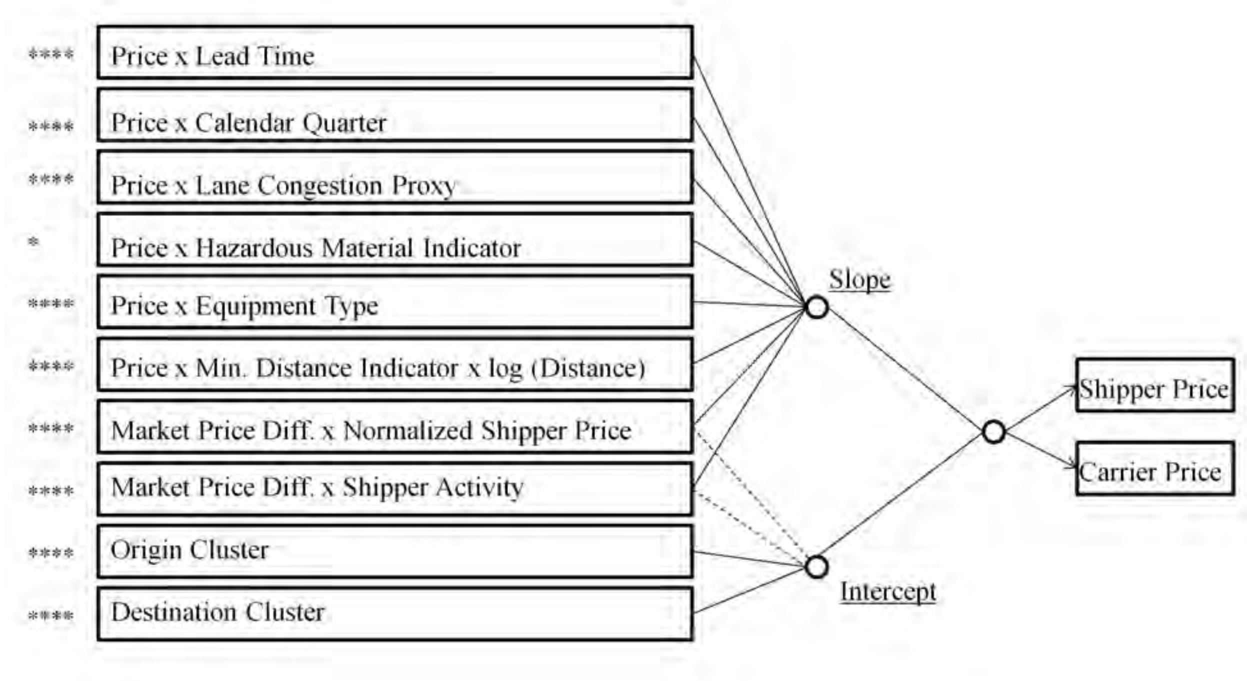

Figure 3. Case study calibration results for shippers and carriers. Significance levels: ‘*’ > 0.20, ‘**’ = [0.20, 0.10), ‘***’ = [0.10, 0.05), and ‘****’ = 0.05 and below.

The asterisks on the far left-hand side in Figure 3 denote the relative statistical significance of each variable according to the t-statistics produced during estimation. Because many of the variables are categorical, meaning that there are multiple levels but only one to which the actual continuous measure is assigned, the most significant t-statistic across the category is used. For example, there are four calendar quarters in a year so the largest t-statistic of all four quarters is used to report significance in Figure 3. The results indicate that for both shippers and carriers, price, origin, and destination are among the most significant attributes in the acceptance/rejection decision. Overall, parameter estimates are intuitive and consistent with expected shipper and carrier behavior.

3.3.2 Profit Optimization Results and Discussion

In this section, we present the results of the profit maximization analysis. We show three primary results: (1) the shipper and (2) carrier prices produced by the framework are reasonable, and (3) also that the estimated profits are higher than historical profits for a similar level of overall price acceptability. For the profit analysis, we use a holdout sample of tens of thousands of shipments taken from the May 2011 to April 2012 time period. The sample was taken over a full calendar year to ensure that there were observations for all of the temporally based variables.

It is important to determine whether or not the shipper prices produced by the model are in line with what would be expected in the real world. In Figure 4 (a) we see the effect that varying the α parameter has on the mean shipper prices produced by the model. (Due to confidentiality we omit many absolute values in the results and cannot reveal true acceptance rates for shippers and carriers. However, the acceptance rates of the presented results are always comparable to the observed acceptance rates.) Because an α value equal to 0.40 produces mean prices and an overall acceptance rate that is close to what is found in the data, we use it as a baseline for all of the following results. In Figure 4 (b), it is clear that the framework produces prices that are distributed similarly to what was historically offered by the 3PL provider. However, the model prices are more tightly centered on the mean of the historical price distribution. By varying the α parameter, the framework is able to produce results that are either farther away from or closer to the historically observed overall acceptance rate. With this parameter, a 3PL provider may decide on an acceptable level of risk relative to the company's historical overall acceptance rate.

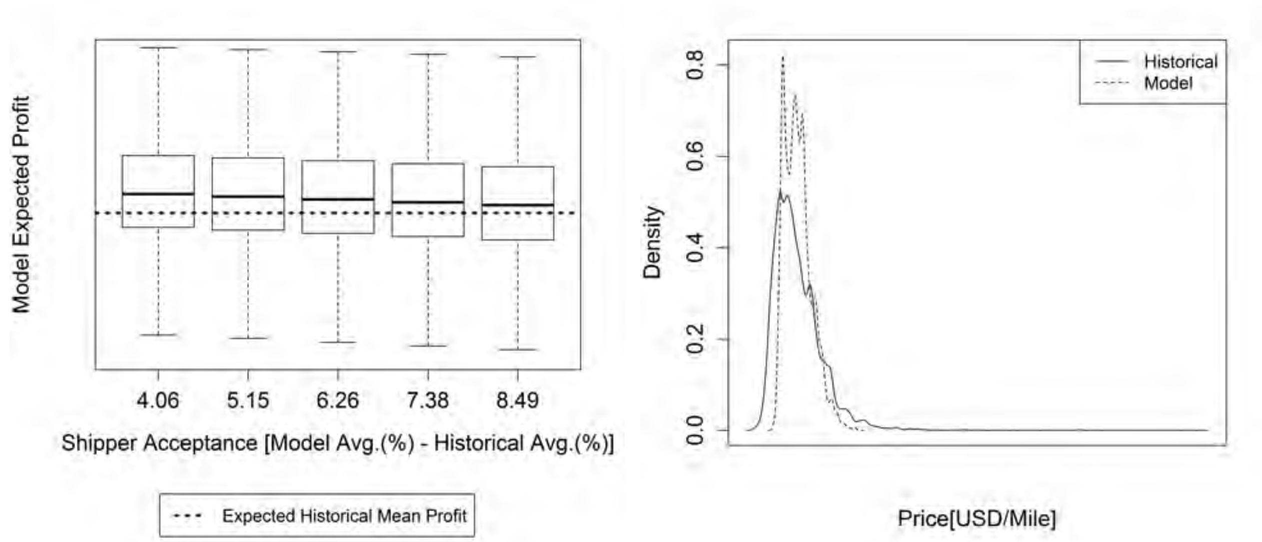

(a) Influence of the acceptance rate on model prices

(b) Density of the true shipper prices and model estimated prices

Figure 4. Comparison of optimal and historical shipper prices.

With regard to carrier prices, in Section 2.2.3 we described the methodology used to develop carrier prices as input to the optimization model. Here, we compare the results of that procedure to a matrix of exogenous average carrier prices on lanes covering the entirety of the contiguous U.S. This matrix of prices was produced by an independent third party. The results in this section use five carriers as its baseline. Five was chosen as the number of carriers because on average that is the amount of carriers the 3PL provider must contact before capacity for a shipment is successfully sourced. In Figures 5 (a) and (b), it is clear that the carrier prices produced by the framework are very similar to average prices on those lanes. Figure 5 (a) is a scatter plot of model and matrix estimated carrier prices. The dashed line labeled “Ideal" is simply a 45-degree line. The closer the scattered points adhere to this line, the closer our estimated prices are to the third-party estimates. That the model results are not far off from the “Ideal" line suggests that the estimated carrier prices are quite accurate. Figure 5 (b) shows this same information but in the form of a histogram. Though the model carrier prices are very close to those of the matrix, they are skewed slightly higher on average.

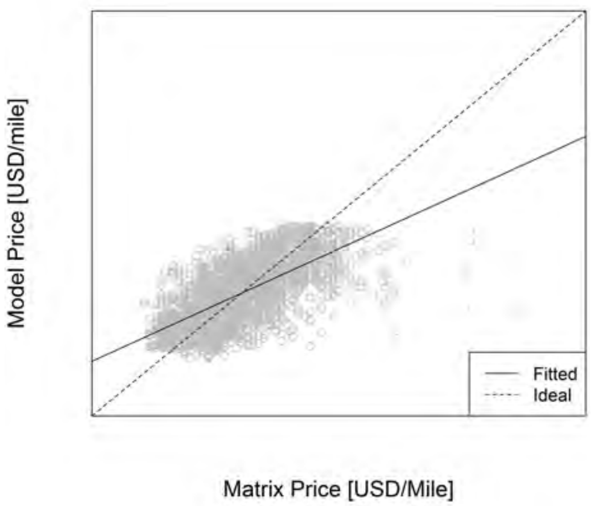

(a) Simulated carrier prices versus matrix carrier prices

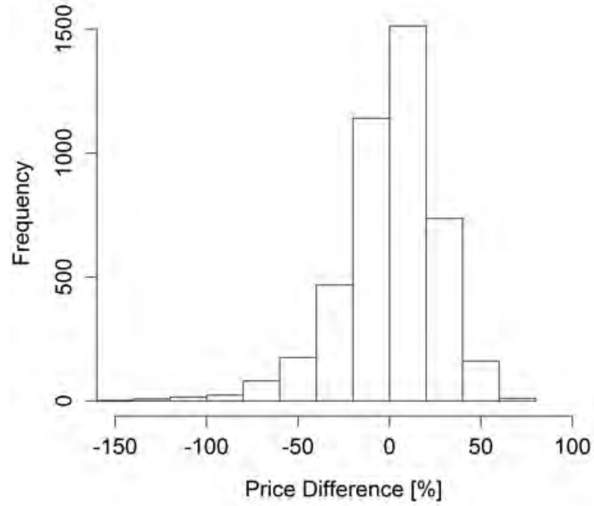

(b) Histogram of the difference between the model and matrix prices (i.e., Matrix price - Model price)

Figure 5. Simulated carrier prices versus matrix carrier prices.

Lastly, we discuss the profit-related results. This analysis has been conducted using a holdout sample of historical shipments, i.e. a sample of shipments that were not used in the estimation of the behavior model parameters. First, in order to gauge how the framework would perform relative to what the 3PL provider currently does, we calculated the difference between the model's “ideal" profit and the 3PL provider historical profit. The ideal profit results assume that shippers always accept the model-generated price - hence the moniker ideal profit. It is calculated as: $Ideal\ Profit\ =\ Optimal\ Shipper\ Price - Historical\ Carrier\ Price.$

Table 1 presents summary statistics of this analysis. Overall, the model produces a larger profit than what was historically observed. However, it also produces slightly more unprofitable shipments and a larger per-mile loss in those unprofitable transactions. The likely reason for this is that those shipments occur on lanes that greatly deviate from the norm in a manner that is difficult to observe, and therefore difficult to capture in the behavioral models; this is especially true when we consider that the data mining process includes a fair amount of spatial aggregation. In reality, the 3PL brokers are probably very familiar with these lanes and are able to adjust their pricing schemes accordingly. Furthermore, over time, and with more observations of such situations, the model could be updated to better recognize such instances.

In Figure 6 (a), the histogram shows the difference in profit as measured by subtracting the historical profit from the ideal profit. The difference is positive when the model suggests a higher profit than what historically occurred and negative when lower. In most cases, the model produces a higher profit for a similar overall acceptance rate.

Next, we analyze the model results from the perspective of a "weighted ideal profit." Essentially, the weighted ideal profit is akin to a weighted average. It takes into account the probability that the shipper may or may not accept the model-generated prices. It is calculated as:
$Weighted\ Ideal\ Profit = \Pr\left(p_S^{Optimal}\right) * Ideal\ Profit.$

We compare this figure to the "weighted historical profit" which is likewise calculated as:

$Weighted\ Historical\ Profit = \Pr\left(p_S^{Historical}\right) * Historical\ Profit,$

where $Historical\ Profit = Historical\ Shipper\ Price - Historical\ Carrier\ Price.$
The probabilities $\Pr\left(p_S^{Optimal}\right)$and $\Pr\left(p_S^{Historical}\right)$ are calculated by inputting the historical and optimal shipper prices into the calibrated logit function, $F(p_S)$. Again, the point of comparing these figures is to gauge to method's ability to produce profitable transactions while taking into account the risk of the shipper not accepting the deal. Figure 6 (b) shows that even when this risk is included, the method is still projected to produce higher profits on average for the 3PL provider. Like the results shown in Figure 6 (a), the overall acceptance rate for the weighted ideal profit is very similar to what was historically observed.

4. Conclusion

This paper presents a modeling framework for the real-time optimal pricing of shipments and sourcing of carrier capacity for a non-asset based 3PL provider in spot markets. The framework was applied to a case study using data from a 3PL provider, demonstrating proof-of-concept. The linking of the core models, discrete-choice and profit maximization, into a framework for 3PL provider pricing and sourcing is a novel application. The framework provides a tool by which a 3PL provider can quickly price shipments and source capacity in spot markets while balancing the risk inherent to the overall process.

Table 1. Summary statistics of the profit analysis. The percentage values show the difference from the historical values.

Profit Measure	**Mean**	**25th Percentile**	**75th Percentile**	**Std. Deviation**
Profit per Mile of Profitable Shipments	8.76%	56.49%	15.39%	-41.99%
Loss per Mile of Unprofitable Shipments	74.76%	88.97%	-70.79%	79.45%
Profit Measure	**Total Change**			
Profitable Shipments	-5.41%			
Unprofitable Shipments	5.41%			
Total Profit	4.90%			

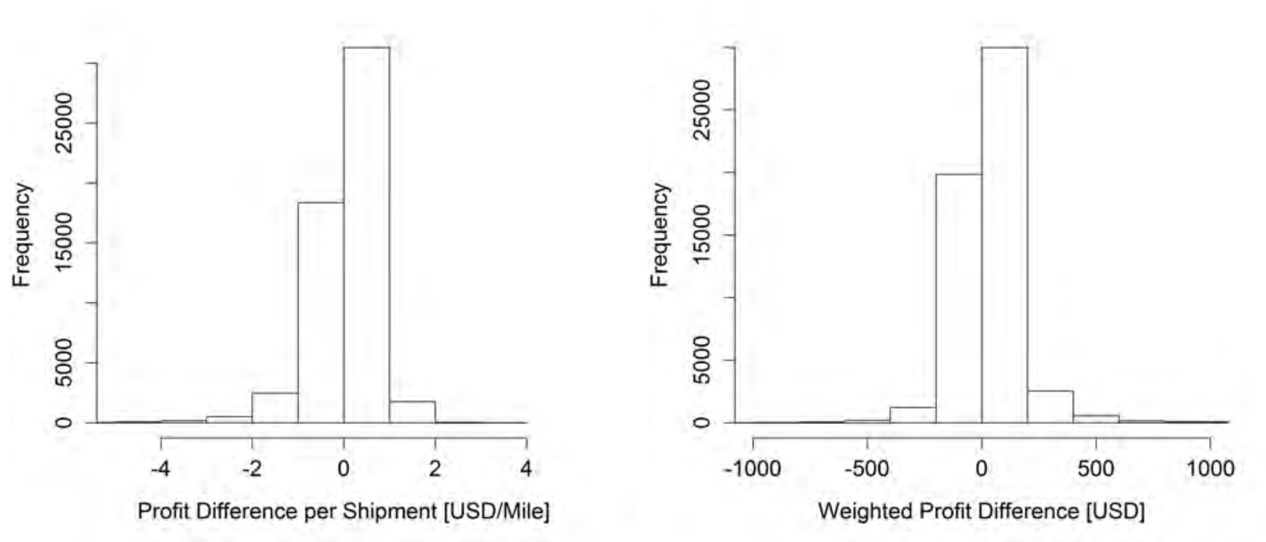

(a) Difference between the model ideal profit and the historical 3PL provider profit (i.e., Ideal profit - Historical profit)

(b) Difference between the weighted ideal profit and the weighted historical 3PL provider profit (i.e., Weighted ideal profit - Weighted historical profit)

Figure 6. Results for the ideal and weighted model profit.

The framework offers a powerful methodology for systematically utilizing transaction data to maximize 3PL profitability while improving the process of matching loads with carriers in the freight marketplace. Though it is of course preferable to use unrestricted choice data, the framework is robust to even choice-restricted data sets, as demonstrated in the case study. The profit maximization model along with the carrier choice set generation procedure balances risk with a 3PL provider's need to quickly source capacity. The carrier ranking procedure provides the 3PL broker with a candidate carrier list that eases the broker's task of identifying good potential carriers on lanes that may be unfamiliar to the broker. Lastly, the data requirement of the framework is not burdensome as 3PL providers regularly capture the type of information needed for the framework in their databases, making our framework easily transferable.

Systems intended to leverage data for better operations and profitability must continually adapt and learn as the marketplace evolves and the key players adjust their behavior. The framework offers the opportunity to learn through the rolling horizon implementation and adaptive process for parameter estimation. Over time, one would expect additional features that reflect changing marketplace conditions and more sophisticated responses from carriers and shippers seeking further competitive advantage. At that time, greater role is envisioned for game theoretic constructs that recognize a richer set of behavioral and strategic responses by these players.

References

1. Broyden, C., 1970. The convergence of single-rank quasi-Newton methods. *Mathematics of Computation* 24 (10), 365-382.
2. Caplice, C., 1996. An optimization based bidding process: a new framework for shipper-carrier relationships. Ph.D. thesis, Massachusetts Institute of Technology, Dept. of Civil and Environmental Engineering.
3. Caplice, C., Sheffi, Y., 2003. Optimization-Based Procurement for Transportation Services. *Journal of Business Logistics* 24 (2), 109-128.
4. Cheng, H. and X. Qi., 2011. The 3PL's optimal transportation pricing problem in dynamic lot sizing models. In: *The 2011 8th International Conference on Service Systems and Service Management*, 1-5.
5. Cosslett, S. R., 1981. Maximum Likelihood Estimator for Choice-Based Samples. *Econometrica*, 49 (5), 1289-1316.
6. Figliozzi, M., 2004. Performance and analysis of spot truck-load procurement markets using sequential auctions. Ph.D. thesis, University of Maryland, College Park.

7. Figliozzi, M., Mahmassani, HS, and Jaillet, P., 2003. Framework for study of carrier strategies in auction-based transportation marketplace. *Transportation Research Record: Journal of the Transportation Research Board* 1854, 162-170.
8. Figliozzi, M., Mahmassani, HS, and Jaillet, P., 2007. Pricing in dynamic vehicle routing problems. *Transportation Science* 41(3), 302-318.
9. Figliozzi, M., Mahmassani, HS, and Jaillet, P., 2008. Repeated auction games and learning dynamics in electronic logistics marketplaces: Complexity, bounded rationality, and regulation through information. In: Helbing (ed.) *Managing Complexity: Insights, Concepts, Applications*, 137-175, Springer Berlin Heidelberg
10. Fletcher, R., 1970. A new approach to variable metric algorithms. *The Computer Journal* 13 (3), 317-322.
11. Goldfarb, D., 1970. A family of variable metric methods derived by variational means. *Mathematics of Computation* 24 (109), 23-26.
12. Hartigan, J., Wong, M., 1979. Algorithm AS 136: A K-means clustering algorithm. *Journal of the Royal Statistical Society. Series C (Applied Statistics)* 28 (1), 100-108.
13. Huang, M., Homem-de Mello, T., Smilowitz, K., Driegert, B., 2011. Supply Chain Broker Operations. *Transportation Research Record: Journal of the Transportation Research Board* 2224, 1-7.
14. Ledyard, J., Olson, M., Porter, D., Swanson, J., Torma, D., 2002. The first use of a combined-value auction for transportation services. *Interfaces* 32 (5), 4-12.
15. Lindsey, C., Frei, A., Alibabai, H., Mahmassani, H.S., Park, Y.W., Klabjan, D., Reed, M., Langheim, G., Keating T., 2013. Modeling Carrier Truckload Freight Rates in Spot Markets. Technical Report available at http://www.klabjan.dynresmanagement.com.
16. McFadden, D., 1984. Econometric analysis of qualitative response models. *Handbook of Econometrics*, 2, 1395-1457.
17. Selviaridis, K., Spring, M., 2007. Third party logistics: a literature review and research agenda. *International Journal of Logistics Management* 18 (1), 125-150.
18. Shanno, D., 1970. Conditioning of quasi-Newton methods for function minimization. *Mathematics of Computation* 24 (111), 647-656.
19. Sheffi, Y., 1990. Third Party Logistics: Present and Future Prospects. *Journal of Business Logistics* 11 (2), 27-39.
20. Sheffi, Y., 2004. Combinatorial auctions in the procurement of transportation services. *Interfaces* 34 (4), 245-252.
21. Song, J., Regan, A., 2003. Combinatorial auctions for transportation service procurement: The carrier perspective. *Transportation Research Record: Journal of the Transportation Research Board* 1833, 40-46.
22. Steinberg, D., Cardell, N., 1992. Estimating logistic regression models when the dependent variable has no variance. *Communications in Statistics – Theory and Methods* 21 (2), 423-450.

Appendix A. Behavior Model Calibration Results

Based on the data, several variables were created in order to capture the effects of shipper, carrier, shipment and lane attributes on the acceptance or rejection of an offer or bid. The variables used in the final calibrated models along with their definitions are included in Table A.1.

Table A.1. Variable definitions.

Variable (Units)	**Description**
Shipment-specific variables	
ppm ($/mi)	Carrier price per mile
rpm ($/mi)	Shipper revenue per mile
Lead Time (Low, High)	Lead time is the time difference in days between the shipment pick-up date and the carrier booking date. It is included as an indicator variable: Low = lead time ≤ 3 days; High = lead time > 3 days
Equipment (-)	Equipment type used for the shipment: Van, Refrigerated, Flatbed (reference category), and Other
dppmC ($/mi)	The difference between the accepted carrier price and the market price per mile
dppmS ($/mi)	The difference between the accepted shipper price and the market price per mile
HazMat	A binary variable indicating whether or not the shipment consists of hazardous materials
Miles (mi)	The total distance of the shipment
MinDist (Short, Long)	An indicator variable describing the distance of the shipment: Short = distance ≤ 300 mi, Long = distance > 300 mi
Quarter (-)	The calendar quarter in which the shipment occurred: 1, 2, 3, or 4
OriginCluster (-)	The cluster in which the shipment originates
DestCluster (-)	The cluster for which the shipment is destined
Lane-specific variables	
Market Price ($/mi)	Average of the median lane ppm and rpm
Mean Lane PPM ($/mi)	Average ppm over a given time period
Mean Lane RPM ($/mi)	Average rpm over a given time period
Lane Volume to Capacity Ratio (Low, Moderate, High)	A categorical variable describing the number of shipments in a lane divided by the no. of carriers in a lane over a given time period: Low (< 33rd percentile), Moderate (33rd - 67th percentile), High (> 67th percentile)

Variable (Units)	**Description**
Carrier-specific variables	
Normalized Mean Carrier PPM (-)	Mean carrier ppm divided by the Mean Lane PPM over a given time period
Carrier Activity Code (Low, MLow, MHigh, High)	A categorical variable describing the number of shipments a carrier has with the 3PL provider: Low (< 25th percentile), MLow = moderately low (25th -50th percentile), MHigh = moderately high (50th -75th percentile), and High (> 75th percentile)
Shipper-specific variables	
Normalized Mean Shipper RPM (-)	Mean shipper rpm divided by the Mean Lane RPM over a given time period
Shipper Activity Code (Low, MLow, MHigh, High)	Categorical variable describing the number of shipments a shipper has with the 3PL provider: Low (< 25th percentile), MLow = moderately low (25th -50th percentile), MHigh = moderately high (50th -75th percentile), and High (> 75th percentile)

The full choice model calibration results for the shippers and carriers are shown in the following tables. Parameter estimates are reported as normalized values using the respective variable's mean as the normalizing value (see Table A.2 and Table A.3).

Table A.2. Shipper acceptance model calibration results.

Variable	**Coef.**	**t-statistic**
Intercept	0.560	8.41
rpm:LeadTime - Low	-0.118	-5.46
rpm:LeadTime - High	-0.104	-4.98
rpm:Quarter2	0.027045	4.01
rpm:Quarter3	-0.00806	-1.14
rpm:Quarter4	0.00162	0.224
rpm:HazMat -Hazardous	0.0382	1.09
rpm:LaneVolumeToCapacityCode - High	0.0471	4.19
rpm:LaneVolumeToCapacityCode - Moderate	0.0276	2.44
rpm:Equipment - Other	-0.0267	-1.702
rpm:Equipment - Refrigerated	-0.0410	-1.407
rpm:Equipment - Van	-0.0146	-2.26
dppmS:NMShipperRPMc - Low	0.0765	5.27
dppmS:NMShipperRPMc - High	0.105	7.53
dppmS:NMShipperRPMc - MHigh	0.0228	0.927
dppmS:NMShipperRPMc - MLow	-0.0793	-2.46
dppmS:ShipperActivityCode - High	-0.0959	-6.07
dppmS:ShipperActivityCode - MHigh	0.0103	1.0496
dppmS:ShipperActivityCode - MLow	0.002625	0.260
rpm:MinDist - Long:log(Miles)	-0.0411	-10.2
rpm:MinDist - Short:log(Miles)	-0.0258	-7.02
OriginCluster2	0.0149	0.243
OriginCluster3	-0.423	-6.00
OriginCluster4	-0.510	-6.60
OriginCluster5	-0.267	-3.31
OriginCluster6	-0.550	-7.66
OriginCluster7	-0.191	-2.92
OriginCluster8	-0.546	-3.18
OriginCluster9	-0.0299	-0.343
OriginCluster10	-0.515	-5.98
OriginCluster11	-0.501	-6.06
OriginCluster12	-0.769	-10.7
OriginCluster13	-0.319	-3.47
OriginCluster14	-0.885	-9.72
OriginCluster15	-0.187	-2.26
OriginCluster16	-1.10	-11.8
OriginCluster17	-0.637	-8.11
OriginCluster18	-0.725	-10.4

Variable	**Coef.**	**t-statistic**
OriginCluster19	-0.318	-4.25
OriginCluster20	-0.0579	-0.700
DestCluster3	-0.0877	-1.204
DestCluster4	0.796	9.68
DestCluster5	0.0998	1.22
DestCluster6	-0.0719	-0.935
DestCluster7	-0.0531	-0.688
DestCluster8	0.712	5.817
DestCluster9	0.0691	0.783
DestCluster10	0.144	1.599
DestCluster11	-0.116	-1.57
DestCluster12	-0.07301	-0.977
DestCluster13	-0.112	-1.40
DestCluster14	0.315	3.54
DestCluster15	0.000323	0.00360
DestCluster16	0.346	3.81
DestCluster17	0.230	2.82
DestCluster18	0.514	6.45
DestCluster19	-0.0293	-0.380
Likelihood Ratio Index	0.0436	

Table A.3. Carrier acceptance (sourcing) model calibration results.

Variable	Coef.	t-statistic
Intercept	-2.424	-27.7
ppm:LeadTime - Low	-0.218	-4.60
ppm:LeadTime - High	-0.207	-4.27
ppm:Quarter2	0.0272	1.89
ppm:Quarter3	0.0216	1.52
ppm:Quarter4	0.0358	2.52
ppm:HazMat - Hazardous	-0.0616	-0.479
ppm:LaneVolumeToCapacityCode - High	-0.0720	-2.72
ppm:LaneVolumeToCapacityCode - Moderate	-0.0751	-3.04
ppm:Equipment - Other	-0.138	-3.81
ppm:Equipment - Refrigerated	0.000265	0.00500
ppm:Equipment - Van	-0.00518	-0.302
dppmC:NMCarrierPPMc - Low	0.109	4.48
dppmC:NMCarrierPPMc - High	0.117	4.62
dppmC:NMCarrierPPMc - MHigh	0.215	4.20
dppmC:NMCarrierPPMc - MLow	0.112	3.70
dppmC:CarrierActivityCode - High	0.0443	1.94
dppmC:CarrierActivityCode - MHigh	0.0390	1.58
dppmC:CarrierActivityCode - MLow	0.0136	0.578
ppm:MinDist -Long:log(Miles)	0.148	18.0
ppm:MinDist -Short:log(Miles)	0.126	15.7
OriginCluster2	-0.176	-2.05
OriginCluster3	0.145	1.48
OriginCluster4	0.353	3.29
OriginCluster5	-0.00511	-0.0457
OriginCluster6	-0.0880	-0.897
OriginCluster7	-0.0607	-0.668
OriginCluster8	0.183	0.776
OriginCluster9	-0.130	-1.08
OriginCluster10	0.0994	0.838
OriginCluster11	0.211	1.88
OriginCluster12	0.360	3.70
OriginCluster13	-0.142	-1.13
OriginCluster14	0.210	1.74
OriginCluster15	-0.141	-1.22
OriginCluster16	0.231	1.88
OriginCluster17	0.219	2.04
OriginCluster18	0.220	2.29
OriginCluster19	0.0326	0.316
OriginCluster20	-0.104	-0.904

Variable	Coef.	t-statistic
DestCluster2	0.0205	0.225
DestCluster3	-0.130	-1.31
DestCluster4	-0.736	-6.30
DestCluster5	-0.362	-3.18
DestCluster6	-0.361	-3.43
DestCluster7	-0.110	-1.04
DestCluster8	-0.327	-1.95
DestCluster9	-0.0586	-0.486
DestCluster10	-0.413	-3.33
DestCluster11	-0.276	-2.69
DestCluster12	-0.141	-1.38
DestCluster13	-0.183	-1.65
DestCluster14	-0.502	-4.20
DestCluster15	-0.0965	-0.791
DestCluster16	-0.538	-4.28
DestCluster17	-0.299	-2.65
DestCluster18	-0.571	-5.16
DestCluster19	-0.178	-1.69
DestCluster20	-0.0654	-0.604
Likelihood Ratio Index	0.0328	