

# Long-Term Relationships in the US Truckload Freight Industry

Adam Harris and Thi Mai Anh Nguyen\*

May 22, 2023

[Click here for the latest version.](#)

This paper provides evidence on the scope and incentive mechanisms of long-term relationships in the US truckload freight industry. In this setting, shippers and carriers engage in repeated interactions under fixed-rate contracts that allow for inefficient opportunism. The main dynamic mechanism involves shippers using the threat of relationship termination to deter carriers from short-term opportunism. This threat and the potential of future rents induce more carrier cooperation. We test this mechanism against likely alternatives and analyze relationship scope for different carrier types. We find that incentive schemes do not exploit the full temporal and spatial scope of relationships.

The importance and ubiquity of informal interfirm relationships is widely recognized. As the economics, management, and sociology literatures have documented, where contracts do not exist or are incomplete, interfirm relationships are governed by goodwill, trust, and reciprocity.<sup>1</sup> A wide range of theoretical work has elucidated various reasons why such informal arrangements might exist, what form they might take, and how they might be sustained.<sup>2</sup> A budding empirical literature studies these relationships. This paper contributes to that empirical literature by studying long-term

---

\*Harris: Department of Economics, Massachusetts Institute of Technology, 50 Memorial Dr, Bldg E52, Cambridge, MA 02142, USA (email: asharris@mit.edu). Nguyen: Department of Economics, Massachusetts Institute of Technology, 50 Memorial Dr, Bldg E52, Cambridge, MA 02142, USA (email: anhng@mit.edu). This research was made possible by the many people in the truckload freight industry who generously shared their time, insights, and data with us. We are especially grateful to Steve Raetz of CH Robinson—whose expertise we refer to throughout the paper—for both the TMC data set, as well as his invaluable feedback on our work. We are also very grateful to Angi Acocella, Chris Caplice, Glenn Ellison, Bob Gibbons, Tom Hubbard, Stephen Morris, Nancy Rose, Tobias Salz, and Mike Whinston for their advice and comments. We also thank all participants in the MIT Industrial Organization Lunch, the MIT Organizational Economics Lunch, the MIT Freight Lab, and the IIOC Rising Stars session for their comments and feedback. We acknowledge the support of the George and Obie Shultz Fund. This material is also based upon work supported by the National Science Foundation Graduate Research Fellowship under Grant No. 1745302. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

<sup>1</sup>Early examples include Macaulay (1963).

<sup>2</sup>For an overview, see Malcomson (2010).

relationships in the US truckload freight industry. In this setting, we ask (1) what is the mechanism governing shipper-carrier relationships, and (2) what is the scope—both temporal and spatial—of that mechanism. The central role of informal long-term relationships, along with the existence of microdata on every interaction within these relationships, makes this setting particularly well-suited for studying these questions.<sup>3</sup>

Sections 1 and 2 of the paper provide a detailed description of the setting and data. The microdata that makes our analysis possible comes from a transportation management system (TMS) used by shippers to manage their relationships with carriers. The data records every interaction within these relationships, the outcomes of the requests for proposals (RFPs) through which the relationships are formed, and shippers’ rankings of carriers, which indicate the status of relationships. The observability of shipper’s play and carriers’ responses enables us to shed light on both the mechanism governing relationships and the scope of that mechanism.

Section 3 then uses the data to establish two key facts about the interactions between shippers and carriers. First, the spot market creates a temptation for carrier deviation; however, carriers do not behave as opportunistically as we might expect if they were playing static best responses. Second, shippers control relationship termination, a power that can potentially be used to punish carrier opportunism. This hypothesized punishment mechanism could explain carriers’ apparent resistance to opportunistic behavior.

Section 4 proposes a parsimonious model consistent with these facts in which a punishment mechanism governs the shipper-carrier relationship: the shipper uses the threat of relationship termination, leveraging the carrier’s value of future rents from the relationship, to incentivize greater carrier cooperation.<sup>4</sup> From the model, we derive testable predictions about qualitative features of the optimal incentive scheme and carriers’ dynamic responses. While the model is intentionally minimal, it serves an important role: guiding our approach to econometric challenges and possible alternative mechanisms in our empirical analysis.

Building on the model, Section 5 returns to our main questions, empirically testing and quantifying (1) the incentive mechanisms and (2) their temporal and spatial scope.

First, to assess the question of *temporal scope*, we examine carrier behavior in the contract’s final weeks. We use mass RFP events as a plausibly exogenous source of relationship termination.

---

<sup>3</sup>About 80% of total industry volume is arranged through long-term shipper-carrier relationships, which, as we will describe in Section 1, are largely informal.

<sup>4</sup>This proposed punishment mechanism is by no means unique to the trucking setting; it is, for instance, common in employer-employee relationships, where the threat of termination deters employee shirking. This threat is stronger when the wage—and thus future rents—is higher. This gives rise to the “efficiency wage” (Shapiro and Stiglitz, 1984), whereby the employer pays an above-market-clearing wage. A similar phenomenon is likely relevant in our setting: the shipper may prefer a higher contract rate, as this incentivizes carrier cooperation via two channels: (1) it increases the carrier’s static payoff from accepting, and (2) it increases the carrier’s future rents from the relationship. While the determination of the contract rate—the formal component of the relationship—and its interaction with informal incentives are fascinating issues, they are beyond the scope of this paper.

We find evidence of endgame effects, with carriers reducing their tendency to accept loads by 10-18 percentage points after learning of the contract period’s imminent end. We argue that these findings indicate that carriers are highly responsive to dynamic incentives before the contract’s final weeks and that the incentive mechanism’s temporal scope is limited to within rather than across contract periods.

Second, to assess the question of *spatial scope*, we quantify the degree to which relationship status on one lane is conditioned on the carrier’s performance on the same lane versus on other lanes.<sup>5</sup> As we would expect carriers’ cost structure to differ based on their size and asset ownership, we naturally estimate shippers’ incentive schemes separately for different carrier types: large asset-based carriers, small asset-based carriers, and brokers. This both provides a more nuanced answer to the question of scope and allows us to speak to the mechanisms at play. We find that large asset-based carriers face a harsh single-lane punishment scheme; a rejection on one lane can reduce the expected duration of the relationship on that lane by up to five loads (from a baseline of 63.5 loads). Brokers, in contrast, face a multi-lane punishment scheme.

Third, we quantify carriers’ responses to the resultant dynamic incentives by estimating how carrier acceptance responds to own-lane volume, an exogenous proxy for the continuation value of the relationship.<sup>6</sup> The ordering of carrier types by their estimated responses to own-lane volume aligns with the strength of the estimated punishment schemes. Large asset-based carriers are most responsive to volume: doubling volume increases their acceptance probability by 6.2pp. As this would not be true for a carrier playing static best-response, we take this as strong additional evidence that, as suggested by our finding of endgame effects, carriers respond to dynamic incentives.

Finally, Section 6 discusses several other mechanisms that may explain observed patterns. The mechanisms we consider include shippers learning about carriers’ characteristics and investments. The fact that demotions tend to be permanent in our data refutes learning as the sole mechanism underlying the observed dynamics of shipper-carrier interactions. However, a combination of learning and network adjustments could explain patterns within relationships involving small asset-based carriers, who show less flexibility in adjusting their network of truck movements than do large asset-based carriers and brokers.

Our paper relates to the empirical literature on long-term informal relationships. While this literature is relatively recent, the last two decades have seen the development of a rich body of empirical evidence on the nature and value of these relationships.<sup>7</sup> Different dynamic mechanisms have been explored, including dynamic enforcement (Brugues, 2023), reputation and learning (Macchiavello and Morjaria, 2015), and adaptation (Barron et al., 2020; Gil et al., 2021). The

---

<sup>5</sup>The term “lane” refers to an origin-destination pair.

<sup>6</sup>This use of volume as a proxy for the expected future relationship value is in keeping with Gil and Marion (2013).

<sup>7</sup>For a review of recent work, see Macchiavello and Morjaria (2023).

value and effects of these mechanisms can be quantified by exploiting exogenous variation in spot rates (Macchiavello and Morjaria, 2015; Blouin and Macchiavello, 2019) or changing prospects of future interactions (Gil and Marion, 2013). Our paper builds on both the conceptual and methodological insights of this literature to study shipper-carrier relationships in the US trucking industry. While most of our empirical evidence points toward dynamic enforcement, we also find suggestive evidence that, depending on carriers’ size and asset ownership, other mechanisms may be at play.

Our paper also relates to the literature on the trucking industry. Earlier papers in this literature include Rose (1985, 1987), Hubbard (2001), Baker and Hubbard (2003, 2004), and Masten (2009). Since these papers, technological improvements have generated rich transaction-level data on shipper-carrier relationships. Such data has been used in the transportation and logistics literature to analyze carriers’ load acceptance, a key measure of performance in the trucking industry.<sup>8</sup> For example, Scott et al. (2017) find that carriers’ tendency to accept offers within relationships is positively correlated with the volume and consistency of timing of these offers. Using the same data set as us, Acocella et al. (2020) find that, when shippers maintain high rates during a market downturn, carriers do not reciprocate with higher acceptance during a later market upturn. To the best of our knowledge, our paper is the first to dissect the dynamic mechanisms underlying shipper-carrier relationships.

Relative to these literatures, we make two main contributions. *First*, we study long-term informal relationships in an important yet understudied industry, the US truckload freight industry. Revenue in this industry was \$700 billion in 2015, equivalent to about 4% of US GDP. Moreover, two aspects of this industry—(i) *fixed-rate contracts* and (ii) *on-path termination*—differentiate it from other settings in which relational contracts have been studied.<sup>9</sup> On the one hand, the lack of flexible monetary transfers prevents us from using workhorse models of long-term informal relationships (MacLeod and Malcomson, 1989; Baker et al., 2002; Levin, 2003), which rely on relational bonuses to support optimal stationary contracts.<sup>10</sup> On the other hand, the fact that relationship termination occurs on-path in our setting allows us to directly estimate the incentive contract, which is nonstationary. Our *second* contribution is that, by exploiting a unique data opportunity, we directly test widely held assumptions in the literature on long-term informal relationships. In particular, we show empirically that relationships do not necessarily exist at the firm-to-firm level as predicted by the multi-market contact literature (Bernheim and Whinston, 1990) and assumed in other empirical studies (Gil et al., 2021). To our knowledge, our paper is the

---

<sup>8</sup>For an excellent review of this work, see Acocella and Caplice (2023).

<sup>9</sup>Rationalizing these unique features of the truckload freight setting is beyond the scope of our paper. We will instead take these features as given, allowing us to focus on other aspects of the incentive contracts.

<sup>10</sup>Similarly, in long-term relationships where one side has full commitment power, a characterization of the optimal dynamic contract typically uses a first-order-condition approach that relies on flexible monetary transfers. For example, Bragues (2023) applies Pavan et al. (2014) to characterize dynamic non-linear contracts with limited enforcement between sellers and buyers in the Ecuadorian manufacturing supply chain.

first to test the assumption of relationship scope.<sup>11</sup> Since the relationships studied previously do not feature on-path termination, it would not be possible to test whether incentive power is pooled across all relationship sub-parts (e.g. products or markets). The fact that we observe key aspects of the relationship—its status, the agent’s performance, the agent’s outside option, and the principal’s termination strategy—at a sub-relationship level (in our case, the lane level) permits us to perform such a test.

# 1 Setting

We begin by describing our setting: the US for-hire truckload freight industry. This is an economically important industry in which informal inter-firm relationships play a central role. We describe the distinguishing features of the industry, as well as the market institutions relevant to our analysis.

## 1.1 The US for-hire truckload freight industry

The freight trucking industry plays a uniquely important role in the US goods economy. In 2015, trucks carried 72% of domestic shipments by value.<sup>12</sup> US trucking firms had revenues of more than \$700 billion in 2015, equivalent to nearly 4% of US GDP in 2015.<sup>13</sup>

Within the freight trucking industry, services are differentiated by the contractual relationships between shippers and carriers, by the size of shipments, and by the equipment required. In this paper, we focus on for-hire truckload carriers supplying dry-van services. We will explain each of these terms in turn: First, a *for-hire* carrier is one who sells his services to various different shippers. This is in contrast to a private-fleet carrier, who is vertically integrated with a single shipper. Second, a *truckload* carrier accepts only large shipments that fill all or nearly all of a trailer. Truckload service is “point-to-point”: A truckload shipment has a single origin and a single destination. While a truckload carrier must plan his network of truck movements efficiently to minimize empty miles, his problem is far simpler than the optimization problem faced by a less-than-truckload carrier, who aggregates smaller shipments to fill the trailer.<sup>14</sup> Finally, a freight truck consists of a tractor unit, which contains a heavy-duty towing engine and a driver cab, and a cargo trailer, which holds goods being hauled by the tractor. Some common trailer classes include

---

<sup>11</sup>Hortaçsu et al. (2021) explores a different, albeit related question: Whether airlines optimally coordinate prices across routes. They find evidence of poor coordination across different organizational teams within an airline, resulting in suboptimal pricing decisions.

<sup>12</sup>See Bureau of Transportation Statistics Freight Facts and Figures 2017.

<sup>13</sup>*American Trucking Trends*, American Trucking Association, 2015.

<sup>14</sup>To avoid confusion, we will, throughout the paper, refer to the shipper using she/her pronouns and the carrier using he/him pronouns.

refrigerated, flatbed and tanker. By far the most common trailer type is the *dry van*, used for hauling boxes or pallets of dry goods not requiring refrigeration. We will focus exclusively on dry van truckload services supplied by for-hire carriers. This is the largest subsegment of the trucking industry and one in which carriers' business model and logistical challenges are easy to understand.

## 1.2 Truckload carriers: Asset-based carriers and brokers

Carriers providing truckload services can be divided into two types: asset-based carriers and brokers. An asset-based carrier owns and operates trucks, which he uses to transport goods. A broker, on the other hand, has no trucks; instead, when a broker accepts a load from a shipper, he subcontracts in a spot arrangement with an asset-based carrier to transport the goods.

Among asset-based carriers, there is enormous heterogeneity in fleet size. On one extreme, the largest US carriers each operate more than 10,000 trucks. On the other extreme, as of December 2020, there were 317,791 registered carriers operating only a single truck.<sup>15</sup> While these tiniest owner-operator carrier will not feature in our analysis, heterogeneity in carrier size—and thus capacity—will play an important role. Our empirical analysis will distinguish between small asset-based carriers (100 trucks or fewer) and large asset-based carriers (more than 100 trucks).

## 1.3 Market institutions

In the US for-hire truckload freight market, shippers and carriers arrange loads through two primary market institutions: a spot market and (largely informal) long-term relationships.

Typically, about 20% of loads are arranged through the spot market.<sup>16</sup> The dominant spot market platform is organized by DAT Solutions; this online “load board” is a simple post-and-search marketplace that facilitates matches between shippers with loads and carriers with trucks.

The remaining 80% of US truckload transactions are arranged through long-term relationships between shippers and carriers. While these relationships are formalized by contracts, the contracts are highly incomplete. A contract defines liability for lost or damaged goods and establishes the rate the shipper will pay the carrier for each load on the lane. However, it imposes few other restrictions on the parties and does not obligate the shipper and carrier to behave cooperatively toward one another.<sup>17</sup> In particular, the contract does not obligate the carrier to accept any loads offered by the shipper under the terms of the contract. If the carrier rejects loads, the contract does not give the shipper any legal recourse.

---

<sup>15</sup>FMCSA, Motor Carrier Management Information System (MCMIS).

<sup>16</sup>[medium.com/@sambokher/segments-of-u-s-trucking-industry-d872b5fca913](https://medium.com/@sambokher/segments-of-u-s-trucking-industry-d872b5fca913)

<sup>17</sup>The fact that these agreements are informal and not legally binding is widely understood in the industry. For instance, Melton Truck Lines Senior Vice President Dan Taylor wrote “The ‘rate agreements’ and ‘load commitments’ for the most part have no contractual obligation or penalties on either party.” (See Taylor (2011).)

The dominance of long-term relationships in this industry suggests that they offer benefits not enjoyed in spot arrangements.<sup>18</sup> Such benefits could take several forms: First, shippers and carriers who interact repeatedly may benefit from a familiarity with each other’s facilities and processes, which can improve functions like loading and payments. Second, arranging loads through a long-term relationship might save on costs associated with searching and haggling in a thin spot market.<sup>19</sup> Such costs are likely non-negligible, as demand for transportation services is dispersed across space and time. Third, and closely related, because spot-market demand on a particular lane at a particular time might be scarce, carriers may prefer the more consistent demand from contracted shippers, which facilitates a stable, cost-effective network of truck movements for the carrier.

## 1.4 Managing relationships: The routing guide

A shipper frequently has contracts with several different carriers on a particular lane. These various carriers, are not, however, equal in status. The shipper explicitly ranks the carriers in a catalog called the *routing guide*. This ranking specifies the order in which carriers are sequentially offered each load that the shipper has on this lane.<sup>20</sup>

To illustrate this sequential offering process (sometimes called a *waterfall*), Table 1 gives an example of an (anonymized) routing guide for a shipper Z on the lane from City X to City Y. When Z has a load at City X that she wants to ship to City Y on a particular date, she first offers the load to the *primary carrier*, in this case, A. If A accepts, then A carries the load and receives \$1230. If A rejects, then the load is offered to B. If B rejects, the load is offered to C, and so on. If the routing guide is exhausted without a carrier accepting, the shipper will typically turn to the spot market to try to find a carrier to accept the load.

---

<sup>18</sup>Note that while long-term relationships are dominant, not every shipper has a long-term relationship on every lane on which she has demand. For instance, if a shipper’s demand on lane  $\ell$  is very low or unpredictable, she will likely not form a long-term relationship for this lane and rather will simply rely on the spot market when demand arises.

<sup>19</sup>The role of transaction costs in driving the tendency towards contractual arrangements is a well established idea. For related studies in the context of the trucking industry, see Hubbard (2001) and Masten (2009).

<sup>20</sup>For a more detailed discussion of the routing guide and related features of truckload operations, see Section 4 of Caplice (2007). As we discuss in the next section, the process of sequentially offering loads is automated by software called a transportation management system (TMS). The TMS that allows each carrier only a short amount of time to respond to an offer. A typical response window might be 45 minutes. This rapidity suggests that the shipper does not have a strategic incentive to rank a carrier higher just because that carrier is in high demand by other shippers; so little time passes between offers that a lower-ranked carrier is unlikely to be “snatched up” by another shipper while higher-ranked carriers are responding to their offers. This means that a shipper’s static best response is to rank the carriers according to her preference over the carriers. Thus, rejections by top-ranked carriers are generally undesirable for the shipper. In Table 1, for instance, the fact that the shipper chose to rank A above B indicates that she prefers paying \$1230 for service from A to paying \$1327 for service from B. Furthermore, the fact that she ranked B above C, despite the fact that C has markedly lower contract rate, suggests that B provides the shipper with superior service in some dimension other than rate (e.g. quality, reliability). More generally, differences in non-rate characteristics may also be important to the shipper.

Table 1: Example routing guide: Shipper Z, lane City X - City Y (on June 1, 2018)

Order	Carrier	Rate	Type
1	A	\$1230	Primary carrier
2	B	\$1327	} Backup carriers
3	C	\$1095	
4	D	\$1450	

The rationale for the shipper maintaining a routing guide with multiple carriers who have the right to reject loads, rather than a single carrier for whom acceptance is obligatory, is that 100% acceptance by a single carrier is unlikely to be efficient. The demand of a shipper over time is random and, therefore, cannot be perfectly predicted by a carrier.<sup>21</sup> This means that when the shipper offers a load on a particular date, the carrier's trucks may be poorly positioned for carrying this load; doing so may be very costly or infeasible.

There are two features of the routing guide that play an important role in the dynamics of the shipper-carrier relationship:

First, the shipper has discretion to alter the ranking of carriers in the routing guide at any time. Indeed, though Table 1 gives the ordering of carriers on June 1, 2018, the routing guide for the same lane two weeks later is substantially different, with these four carriers ordered C, A, B, D. Why might such a change occur? The shipper might use the power to reorder the routing guide strategically to incentivize carrier cooperation. If a carrier (e.g., Carrier A) were behaving opportunistically, rejecting contract loads in favor of taking higher-paying loads in the spot market, the shipper could punish the carrier by demoting him to a lower position in the routing guide. Being demoted diminishes the carrier's future rents from the relationship, as he will now receive fewer offers on this lane. This possibility of punishment via reorganization of the routing guide is the mechanism at the heart of our model and empirical analysis.

Second, at the end of a contract period, the shipper holds a request for proposals (RFP) to determine the set of carriers, their rates, and their initial positions in the new routing guide. In an RFP, a shipper need not award the primary position to the lowest-bidding carrier; non-rate characteristics can be taken into account. While this is intuitively similar to a scoring auction, the way these RFPs are carried out in practice is far more complicated than the formal auctions that have been studied theoretically and empirically in a wide range of economic settings. After a shipper receives carriers' initial bids, multiple rounds of negotiation between the shipper and the various carriers jointly determine carriers' final routing guide positions and rates.

<sup>21</sup>While the timing of loads is random, the demand of a shipper is typically more consistent than the demand of a single consumer in some other transportation industries, e.g. the taxi or ride-hail industry.



## 2 Data

We use transaction-level data from the transportation management system software used by shippers to manage their relationships with carriers. The data records every interaction within these relationships. To proxy for carriers' outside option, we use a measure of the going rate for freight services in the spot market from DAT, the gold-standard provider of such spot market data.

### 2.1 Shipper-carrier microdata

Our analysis is made possible by the fact that shippers use a transportation management system (TMS) to manage their relationships with carriers and to automate the waterfall of offers. The shipper enters carriers' rates and ranks into the TMS, and then, for each load, prompts the TMS to sequentially send electronic offers to the carriers. For each load sent through the TMS, the software records the details of the load, all offers that are made, and whether each is accepted or rejected. These records for one particular TMS software provider, called TMC, are the source of our microdata.<sup>22</sup>

The microdata covers the period from September 2015 through August 2019. In all, the data set includes 1,074,172 loads and 2,130,125 offers. 71% of loads are accepted by the first carrier to which they are offered. All loads in the data set have a haul distance of at least 250 miles.<sup>23</sup> The mean distance is 692 miles with a standard deviation of 440 miles. The average per-mile contract rate is \$1.85 with a standard deviation of \$0.51.

The time between consecutive RFPs on a lane is, on average, 322 days, with an average of 83 loads being offered during this time. The carrier who wins an RFP on a lane is demoted from primary status before the next auction is held 21% of the time.

**Shipper-carrier relationships and networks** The microdata includes 40 shippers with at least 500 loads. Among these shippers, the median shipper has 8,094 loads with, on average, 192 active lanes and 53 active carriers each year.

On many lanes, shippers and carriers interact infrequently. For example, the median lane of the median shipper has only one load per month. However, among the top 10% of lanes for the median shipper, each has, on average, a load for every four days. Such variation in frequency of interactions will be important for our test of carriers' response to dynamic incentives.

Multilane interactions between a shipper and a carrier are also common. The top five and top ten carriers of the median shipper deliver, respectively, 58% and 73% of her loads. Relatedly, it is

---

<sup>22</sup>TMC is a division of CH Robinson, a third-party logistics firm.

<sup>23</sup>For shorter-distance hauls, the prevailing market institutions are somewhat different. These loads are therefore excluded from our analysis.

common for a carrier to serve as a shipper’s primary carrier on multiple lanes. For example, the top five carriers of the median shipper hold primary status on an average of 21 lanes each. There is thus significant potential for strategic exploitation of multilane interactions: a shipper might condition a carrier’s primary position on one lane on his behavior on another lane. This question of the scope of the incentive mechanism—whether shippers exploit multilane interactions to create cooperative incentives—is one of the key questions this paper addresses.

Shipper-carrier relationships also differ meaningfully across carrier types. As shown in Table 6 in Appendix C, a broker—as compared with an asset-based carrier—tends to serve as the primary carrier on more lanes for a given shippers, but these lanes tend to have lower volume. This suggests greater incentive for and greater scope for pooling incentives across lanes in shippers’ relationships with brokers, something that will be explored in Section 5.2.

## 2.2 Spot rate data

We will use data on the average rate for truckload services in the spot market to capture the relevant outside option—the alternative opportunities available to shippers and carriers outside of their long-term relationships. This data comes from DAT Solutions, the leading provider of data on truckload spot markets. For our sample period, the data set gives us seven-day trailing average spot rates for a set of narrowly-defined lanes that cover the continental United States.<sup>24</sup>

Across all lanes and dates, the overall mean spot rate per mile is \$1.68 with a standard deviation of \$0.60. The first quartile, the median, and the third quartile are \$1.26, \$1.53, and \$1.93, respectively. A notable feature of the data is persistent differences in rates across lanes; a regression of spot rates on a set of lane fixed effects has an  $R^2$  of 0.78, indicating that across-lane differences are large relative to within-lane variation. In later empirical analysis, we pool observations across lanes for the purpose of estimating the strategies of shippers and carriers. To make for appropriate comparisons across lanes, we will use residualized, rather than raw, spot rates, partialling out lane fixed effects. For the time series of average monthly spot rates over our sample period, see Figure 1 in the next section.

## 3 Two Key Facts

In this section, we use our shipper-carrier microdata, together with the data on spot rates, to establish two key facts which suggest that the shipper-carrier relationship can be thought of as a repeated principal-agent game in which an incentive mechanism deters carrier opportunism.

---


<sup>24</sup>Each lane is defined by a pair of key market areas (or KMAs). The continental US is partitioned into 135 KMAs, so there are  $135^2$  KMA-to-KMA lanes.

### 3.1 Fact 1: Temporary spot-contract rate differences create temptation for carriers

We begin by arguing that when spot rates are higher than contract rates, carriers are tempted by short-term opportunism. As such opportunism is detrimental to the shipper, a moral hazard problem exists within the shipper-carrier relationship.

Figure 1 illustrates the potential for carrier opportunism by depicting two key aggregate trends in our relationship microdata and aggregate data on spot rates. First, there are periods in which spot rates are significantly higher than contract rates. Second, these periods coincide with a high proportion of rejections by carriers.

Figure 1: Aggregate trends: National averages of rejection, contract, and spot rates



*Notes:* The monthly rejection rate is constructed from the TMS microdata as the fraction of loads rejected by the first carrier in the routing guide. The average monthly contract rate is constructed from the TMS microdata. The average monthly spot rate is constructed from the DAT data, on the same set of lanes covered by the TMS microdata. The rejection rate, contract rate, and spot rate are all volume-weighted averages.

Figure 1 shows considerably greater intertemporal variation in spot rates than in contract rates. While spot rates were generally lower than contract rates in the first two years of our sample, an aggregate demand shock in late 2017 and early 2018 resulted in a sharp increase in spot rates.<sup>25</sup>

These spot market premia create the potential for short-term opportunism. Recall that a carrier in a long-term relationship always has the option to reject loads offered to him by the shipper.

<sup>25</sup>Contemporary articles from various trade publications (including Transportation Topics and FreightWaves.com) describe the high spot rates of the 2017-2018 period as being driven by increased spending on e-commerce, booming US industrial production, and the December 2017 corporate tax cut. Some sources also cite various supply factors, including the December 2017 introduction of a rule requiring for-hire trucks to be equipped with electronic logging devices (ELDs), though these supply factors seem to be considered of secondary importance. See, for instance, [www.freightwaves.com/news/market-insight/forecasting-2019](http://www.freightwaves.com/news/market-insight/forecasting-2019).

Thus, when the spot rate exceeds the contract rate, the carrier may reject contract loads and instead opt to provide service in the spot market. Figure 1 shows evidence consistent with this hypothesis: The period of high relative spot rates coincides with a large increase in the proportion of offers rejected by primary carriers. This observation strongly suggests that the spot market represents a key outside option for carriers.

Such opportunism by the carrier presents a moral hazard problem. The fact that long-term relationships exist in the first place—rather than all transactions being arranged through the spot market—suggests that there is relationship surplus that would be foregone were the carrier to opportunistically choose to service the spot market.<sup>26</sup> Furthermore, the shipper has imperfect monitoring: the shipper cannot distinguish between an inefficient opportunistic rejection and an efficient rejection resulting from the carrier’s current cost of service being high.

Yet Figure 1 also gives us reason to believe that some mechanism exists to alleviate the moral hazard problem. When spot rates peak in January 2018, they are on average 20% higher than contract rates. Despite this strong incentive for carriers to reject loads, the majority of loads are still accepted by primary carriers in this month. That many carriers are willing to forgo significant short-term profits suggests that their opportunistic tendencies are restrained by some other force.

One such force could be an incentive scheme in which the promise of future rents helps alleviate the carrier’s short-term opportunism. Necessary conditions for such an incentive scheme to be effective are that (i) the shipper has the power to deny the carrier future rents if he behaves opportunistically and (ii) the carrier’s future rents from the relationship are sufficiently large. The next subsection establishes the former; the latter is established in Sections 5.1 and 5.3.

Before addressing the role of shippers in the next subsection, we first address whether, just as carriers behave opportunistically when spot rates are high, shippers also behave opportunistically when spot rates are low. That is, when spot rates are low, do shippers skip over the routing guide and go directly to the spot market? To answer this question, we estimate the following regression:

$$\log(\text{Volume}_{scm}^{\ell}) = \beta_0^{\text{shipper}} + \beta_1^{\text{shipper}} \log(\text{Spot}_m^{\ell}) + \gamma_{sc}^{\ell} + \epsilon_{scm}^{\ell} \quad (1)$$

where  $\text{Volume}_{scm}^{\ell}$  is the number of offers that shipper  $s$  sent to primary carrier  $c$  for service on lane  $\ell$  within month  $m$ , and  $\text{Spot}_m^{\ell}$  is the average spot rate on lane  $\ell$  in month  $m$ . We include shipper-carrier-lane fixed effects  $\gamma_{sc}^{\ell}$  to absorb differences in contract rates, spot rates, and volumes across relationships. If shippers tend to skip the routing guide and go directly to the spot market when spot rates are low, we would expect  $\beta_1^{\text{shipper}} > 0$ . In contrast, we find that  $\hat{\beta}_1^{\text{shipper}} = -0.0363$  with a standard error of 0.0119. We thus conclude that unlike carriers, shippers appear non-responsive to spot temptation.<sup>27</sup>

<sup>26</sup>See Section 1.3 for a discussion of possible sources of this surplus.


<sup>27</sup>Consistent with this finding, industry experts tell us that it is “very rare” for a shipper to go directly to the

### 3.2 Fact 2: Shippers control relationship termination

We next use our shipper-carrier microdata to show that shippers control relationship termination and provide suggestive evidence on the form of shippers' termination strategies.

Figure 2 presents an example of a lane history that motivates the way we think about the shipper's decisions. Recall that the chief decision faced by the shipper is that of when and how to change the routing guide. Such changes can be made at any time. Some are the result of RFPs, while others take place within the contract period, i.e., in the time between RFPs. Our analysis will focus on the latter and, in particular, on those changes that replace one primary carrier with another. We refer to such a change as a *demotion* of the current primary carrier.

Figure 2: An example: Offers for Shipper X, City Y - City Z



*Notes:* Each point represents an offer: circles represent offers that are accepted while crosses represent offers that are rejected. Each carrier is indicated by a different color. The dotted black line indicates the rate of the primary carrier at each point in time. Carrier 1, Carrier 2 and Carrier 3 each serve as primary carrier for this shipper and lane for some subset of the period from October 2017 to May 2019. The offer data comes from the TMS microdata. The average monthly spot rate on the same lane is constructed from the DAT data set.

In the example in Figure 2, Carrier 1 initially holds primary status and accepts most of the spot market before the routing guide when spot rates are low. Perhaps because shipping represents only a small component of the operations of shippers, who are usually non-transportation firms (e.g. manufacturers or retailers), taking advantage of short-term opportunities to reduce shipping costs is not a priority. Shippers allow day-to-day shipping decisions to be automated by the TMS and make strategic decisions only on a medium-term basis.

offers he receives. Around early October 2017, the shipper holds an RFP for this lane; Carrier 1 retains his primary status and gets a rate increase of about 5 cents per mile. However, over the next three months, a period of high spot rates, Carrier 1 rejects many of the loads offered to him. In January 2018, Carrier 1 is demoted from primary status and replaced by Carrier 2. Over the next five months, Carrier 2 rejects most of the loads offered to him. Ultimately, he too is demoted in favor of Carrier 3, who maintains primary status for the rest of the sample period.

This figure, which illustrates patterns that are common to many lanes, motivates two key conclusions about shipper-carrier relationships:

First, while shippers have almost unlimited discretion in what kind of routing guide changes they make, in practice, they do not switch primary carriers frequently; rather, a shipper maintains a primary carrier for a time before ultimately—and usually permanently—demoting that primary carrier.<sup>28</sup> From this observation, it seems appropriate to think of the shipper-carrier relationship in terms of the following kind of principal-agent model: the shipper controls relationship termination and, at each point in time, decides between continuation and (permanent) termination.

Second, a clear pattern on this lane is that a series of rejections by the primary carrier often is followed by a demotion. This pattern is documented more systematically in Section 5.2. This evidence is consistent with an incentive mechanism where the shipper generates dynamic incentives for the carrier by conditioning relationship continuation on acceptance.

The evidence in this subsection indicates that shippers have the power to terminate relationships. Whether the threat of termination is effective in deterring carrier opportunism will be addressed in Sections 5.1 and 5.3, both of which show strong evidence of carriers responding to dynamic incentives.

## 4 A model of the incentive contract

In this section, we develop a model of long-term shipper-carrier relationships that will serve as a theoretical framework for our empirical analysis. For tractability, we focus on the relationship between a shipper and a primary carrier, abstracting away from the existence of backup carriers.<sup>29</sup>

---

<sup>28</sup>See Table 4. We find that in more than 95% of instances where a carrier is demoted from primary status, he never regains primary status on the lane in our sample period (2015-2019). While there is a truncation issue here (a demoted carrier may regain primary status after the end of our sample period), it nevertheless seems clear that demotions are typically permanent. While it is possible that shippers are actually employing reward-punishment cycles like those described by Green and Porter (1984) or Li and Matouschek (2013), the cycles would have to be very long. We observe four years of data—quite a long period of time when one considers that the typical time between consecutive loads on a lane is on the order of a few days—and yet we almost never observe a demoted carrier regaining primary status. Hereafter, we assume that demotion is permanent and use the terms demotion and termination interchangeably.

<sup>29</sup>A model that captured both primary and backup carriers and provided insight as to how a shipper manages relationships with these various carriers would be reminiscent of Board (2011) and Andrews and Barron (2016). While interesting, this issue is beyond the scope of this paper.

Features of our model are motivated by the two facts established in Section 3:

*Fact 1.* Temporary spot-contract rate differences create deviating temptation for carriers,

*Fact 2.* Shippers control relationship termination and can use this power to generate an incentive scheme.

The model also generates predictions that guide our empirical analysis.

**Relationship parameters** A tuple  $(\psi, \eta_1, \eta_2, p, \delta, F, G)$  summarizes the key characteristics of a relationship. Here,  $\psi$  is the relationship-specific gain to the shipper from transacting with the carrier;  $\eta = \eta_1 + \eta_2$  is the relationship-specific gain to the carrier from transacting with the shipper. Some of these gains are publicly observed ( $\eta_1$ ), such as the consistency of timing of shippers' requests, which helps the carrier's planning. Some are privately observed  $(\psi, \eta_2)$ , such as the quality of on-road communication or efficiency of loading and docking. In addition,  $p$  is the contract rate;  $\delta \in (0, 1)$  is the discount factor, measuring the frequency of shipments.<sup>30</sup>  $F$  is the distribution of the carrier's cost of servicing a shipment on the contracted lane, and  $G$  is the distribution of the spot rate on that lane.

Let  $\tilde{p}_t$  and  $c_t$  denote, respectively, the spot rate and the carrier's cost draw in period  $t$ . The shipper's period- $t$  payoff is  $u_t = \psi - p$  if she is served by the contracted carrier and  $u_t = -\tilde{p}_t$  if she is served by the spot market. The carrier's period- $t$  payoff is  $v_t = \eta + p - c_t$  when serving the contracted shipper and  $v_t = \tilde{p}_t - c_t$  when serving the spot market (on the same lane). If the carrier chooses to remain idle (or serve a different lane), he gets zero payoff in that period. Thus,  $c_t$  captures the opportunity cost of servicing the contracted lane in period  $t$ ; the cost distribution  $F$  captures both the contracted lane's average alignment with the rest of the carrier's network and the day-to-day variation of such alignment.

When  $\psi + \eta > 0$ , shipments fulfilled within relationships generate surplus over spot transactions. In this case, it is never efficient for the carrier to reject the shipper's offer in order to serve the spot market on the same lane. However, requiring the carrier to always accept the shipper's load is also not efficient, since the carrier's (opportunity) cost in some periods might be very high.<sup>31</sup> The inability of the shipper to distinguish between rejections due to high cost draws and rejections due to high spot rates represents a potential source of inefficiency in this setting, one that the shipper may hope to alleviate using the threat of relationship termination.

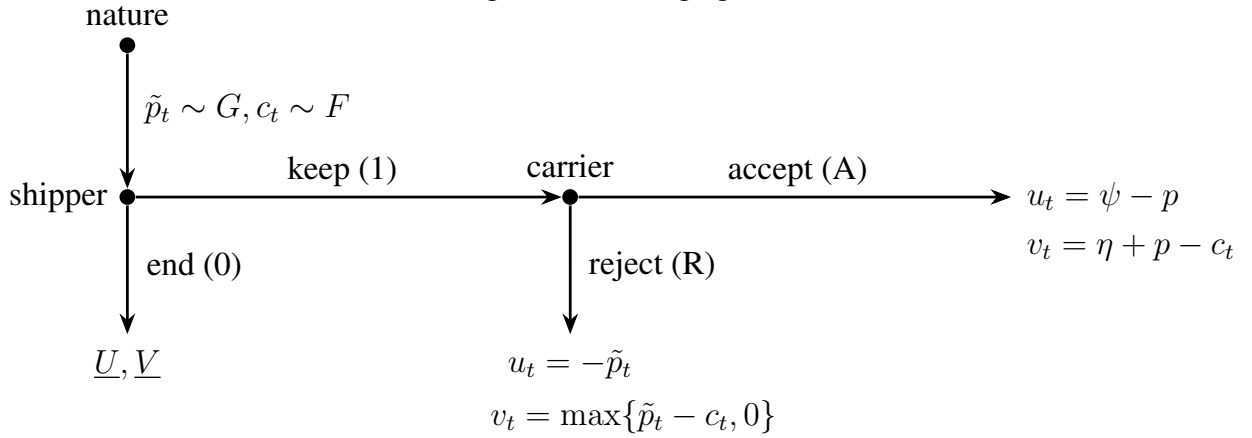
**Timing and informational assumptions** In period  $t = 0$ , the shipper holds an RFP to select a primary carrier. This RFP also reveals to both the shipper and carrier the characteristics

<sup>30</sup>An underlying assumption is that the shipper's demand is perfectly inelastic with respect to spot rates. For evidence that shippers do not reduce load offers within the routing guide when spot rates are low, see Section 3.1.

<sup>31</sup>Specifically, when  $c_t > (\psi + \eta) + \tilde{p}_t$ , the efficient outcome is that the carrier remains idle or serves a different lane.

$(\psi, \eta_1, \eta_2, p, \delta, F, G)$  of their relationships.<sup>32</sup> From period 1 onward, the shipper and the primary carrier interact repeatedly, with spot rates and costs being drawn independently and identically over time.<sup>33</sup> The stage game is summarized in Figure 3. In each period  $t$ , a spot rate  $\tilde{p}_t$  is drawn from  $G$  and publicly observed, and a cost draw  $c_t$  is drawn from  $F$  and privately observed by the carrier. The shipper then decides whether to “keep” the carrier as the primary carrier or “end” their relationship. If the relationship is maintained, the carrier chooses whether to accept (A) or reject (R) the shipper’s load in that period. If he rejects, then he can either serve the spot market or remain idle. Otherwise, if the relationship ends, both sides resort to the spot market for future transactions; the shipper gets expected payoff of  $\underline{U}$  and the carrier gets expected payoff of  $\underline{V}$ .

Figure 3: The stage game



To generate clean predictions, we focus on the simplest class of shipper’s incentive schemes, those that condition only on the carrier’s decision in the last period. Denote such an incentive scheme by  $\sigma_0 : \{A, R\} \rightarrow [0, 1]$ , where  $\sigma_0(d_{t-1})$  is the probability that the shipper maintains the relationship following decision  $d_{t-1}$  of the carrier in the last period. We interpret  $\sigma_0(A)$  as the level of rewards following cooperation and  $1 - \sigma_0(R)$  the level of punishment following noncooperation. We examine the carrier’s optimal response to such scheme. Since the shipper conditions only on the last period’s decision, we can focus on the carrier’s stationary play,  $\sigma_1 : \text{supp}(G) \times \text{supp}(F) \rightarrow [0, 1]$ , where  $\sigma_1(\tilde{p}_t, c_t)$  is the probability that the carrier accepts the offered load given spot rate  $\tilde{p}_t$  and cost draw  $c_t$ .

**Model predictions** Next, we derive testable predictions on the carriers’ optimal stationary play and shipper’s incentive scheme.

<sup>32</sup>See Harris and Nguyen (2023) for modeling assumptions under which relationship characteristics are revealed via the RFP.

<sup>33</sup>Realistically, spot rates exhibit autocorrelation (Figure 1). With autocorrelation, we would expect an increase in the spot rate to have several different effects, as it would change the continuation value of both the shipper and the carrier in addition to increasing contemporaneous temptation for the carrier. Harris and Nguyen (2023) present a model in which spot rates are AR(1) and the shipper keeps track of a rejection index summarizing all past rejections. The qualitative properties of our predictions extend to this more general model.



**Proposition 1.** (Carrier acceptance) Suppose that  $\eta + p \in \text{Supp}(G)$ .<sup>34</sup> The carrier's optimal response takes a threshold form: accept if and only if  $\bar{p} \geq \max\{\tilde{p}_t, c_t\}$ , where

$$\bar{p} = \eta + p + \frac{\delta}{1 - \delta}(V(A) - V(R)),$$

and  $V(d_t)$  is the carrier's expected payoff following  $d_t \in \{A, R\}$ . Moreover, the following comparative statics hold:

i) If  $\sigma_0(A) > \sigma_0(R)$ , then  $\bar{p} > \eta + p$ . That is, the carrier accepts more often than their static best response.

ii) For a fixed incentive scheme with  $\sigma_0(A) > \sigma_0(R)$  and unobserved characteristics  $(\psi, \eta_2, F)$ ,

$$\frac{\partial \bar{p}}{\partial \delta} \geq 0, \quad \frac{\partial \bar{p}}{\partial \eta} \geq 1, \quad \text{and} \quad \frac{\partial^2 \bar{p}}{\partial \delta \partial \eta} \geq 0.$$

iii) For fixed carriers' parameters  $(\eta_1, \eta_2, F)$  and  $\delta$ , and for every  $\sigma_0(A) > \sigma_0(R)$ ,

$$\frac{\partial \bar{p}}{\partial \sigma_0(A)} \geq 0 \quad \text{and} \quad \frac{\partial \bar{p}}{\partial \sigma_0(R)} \leq 0.$$

*Proof.* See Appendix A.1. □

Intuitively, the shipper can use her control over relationship continuation as an incentive scheme to induce dynamic incentives for the carrier to accept more loads. The strength of such dynamic incentives depends on both the value of current and future loads to the carrier, as well as the levels of rewards and punishments induced by the shipper's incentive scheme.

**Proposition 2.** (Shipper's single-lane incentive scheme) Suppose that shipper's match-specific gain is sufficiently large,  $\psi \geq p - \mathbf{E}[\tilde{p}_t | \tilde{p}_t \leq p]$ . The optimal single-lane incentive scheme for the shipper satisfies

i) (Maximum rewards)  $\sigma_0^*(A) = 1$  for any parameter values of  $(\psi, \eta_1, \eta_2, \delta, F, G)$ ,

ii) (Soft punishment)  $\sigma_0^*(R) \in (0, 1)$  for some parameter values of  $(\psi, \eta_1, \eta_2, p, \delta, F, G)$ .

*Proof.* See Appendix A.2.1. □

The incentive scheme affects the shipper's payoffs via two channels: the effect on the carrier's acceptance probability (the incentive-inducing effect) and the direct effect on the probability of

---

<sup>34</sup>That is, there are periods in which spot transactions are more attractive than the contracted offer and periods in which the contracted offer is more attractive than spot transactions.

ending the relationship (the regime-switching effect). Since the shipper faces no tradeoff between these two effects when deciding on the reward, she should choose the maximum reward (guaranteed relationship continuation). In contrast, harsher punishment increases the carrier's acceptance probability but also increases the likelihood of relationship termination when the carrier rejects. It is possible that this tradeoff is not resolved by extreme punishment, but rather by soft punishment.

While Proposition 2 assumes that the shipper's incentive scheme operates at the lane level, Example 1 considers the possibility of a broader scope, operating across multiple lanes within the shipper-carrier relationship. Importantly, part (ii) of this example shows that using a simple scheme to pool incentives across heterogeneous lanes can backfire, making the shipper worse off than if she just used the optimal single-lane incentive schemes.

**Example 1.** *Suppose that the shipper and the carrier interact on two lanes  $\ell = 1, 2$ , each characterized by  $(\psi, \eta_1^\ell, \eta_2^\ell, p^\ell, F, G)$ . Let  $F \sim \alpha U(0, 1) + (1 - \alpha)\delta_K$  for some large  $K$ . That is, a cost draw is with probability  $\alpha$  distributed as a standard uniform random variable and with probability  $(1 - \alpha)$  equal to some  $K \gg 1$ . Let  $G \sim U(0, 1)$ ,  $\alpha = 0.75$ ,  $\delta = 0.9$ ,  $\psi = 0.3$ , and  $p = 0.6$ . Focus on the class of multi-lane incentive schemes that map the average multi-lane rejections  $\frac{1}{2} \sum_{\ell=1}^2 \mathbf{1}\{d_{t-1} = R\}$  of the carrier in the last period to a probability of the shipper keeping the relationship (on both lanes). Denote by  $\hat{\sigma}_0$  the shipper's optimal multi-lane incentive scheme (within this class) and by  $\sigma_0^\ell : \{A, R\} \rightarrow [0, 1]$  the shipper's optimal single-lane incentive scheme for  $\ell = 1, 2$ .*

i) *If  $\eta_1^1 + \eta_2^1 + p^1 = \eta_1^2 + \eta_2^2 + p^2 = 0.65$ , the shipper is better off using  $\hat{\sigma}_0$  than  $(\sigma_0^1, \sigma_0^2)$ .*

ii) *If  $\eta_1^1 + \eta_2^1 + p^1 = 0.65$  and  $\eta_1^2 + \eta_2^2 + p^2 = 0.8$ , the shipper is worse off using  $\hat{\sigma}_0$  than  $(\sigma_0^1, \sigma_0^2)$ .*

*Proof.* See Appendix A.2.2. □

**Empirical challenges** In testing the predictions generated by this model, the econometrician observes a component of carrier's gain ( $\eta_1$ ), the contract rate ( $p$ ), the frequency of interactions ( $\delta$ ), the realized spot rate ( $\tilde{p}_t$ ), and the carrier's decision ( $d_t$ ) in each period; however, other relationship-specific characteristics are unobserved. Our empirical approaches focus on addressing endogeneity issues arising from the unobservability of relationship-specific characteristics ( $\psi, \eta_2$ ) and carrier's cost distribution  $F$ . We will focus on addressing three channels: First, these unobservable characteristics affect the carrier's decisions ( $d_t$ ). Second, they affect the contract rate ( $p$ ) via the RFP process. Third, they potentially correlate with the frequency of interactions ( $\delta$ ). The last of these objects plays an important role in our empirical analysis, serving as a potential shifter for the continuation value of the relationship.

## 5 Empirical Evidence

In this section, we present empirical evidence on two key questions. First, what is the scope—both spatial and temporal—of the incentive mechanisms that govern shipper-carrier relationships? And, second, how do carriers respond to the dynamic incentives created by these mechanisms?

Our analysis of these questions proceeds in three subsections. In Section 5.1, we begin by presenting empirical evidence on the temporal scope of the relationship using carrier acceptance decisions in the weeks preceding the end of the contract period. This exercise also provides preliminary evidence that carriers respond to dynamic incentives. Next, in Section 5.2, we present empirical evidence on the scope of the relationship in the spatial dimension by estimating shippers' demotion strategies. Finally, in Section 5.3, we provide further evidence that carriers respond to dynamic incentives using carrier behavior throughout the relationship. In the latter two subsections, our analysis accounts for an additional dimension of carrier heterogeneity, analyzing behavior separately for three different types of carriers—large asset-based carriers, small asset-based carriers, and brokers.<sup>35</sup>

### 5.1 Empirical Evidence: Carrier behavior at the end of the contract period

This subsection presents evidence that (i) carriers respond to dynamic incentives and (ii) the scope of the incentive scheme carriers face is within, not across, contract periods.

We showed in Proposition 1(i) that a carrier facing dynamic incentives has  $\bar{p} > p + \eta$ , meaning that he is *ceteris paribus* more likely to accept load offers than a carrier with only static incentives. If an exogenous event were to unexpectedly eliminate the carrier's dynamic incentive, then we would expect a decrease in his probability of acceptance. In our data, *mass RFP events* act as a natural experiment which results in such a shock to dynamic incentives.

While we have established that a relationship sometimes ends because the shipper demotes the carrier, it may also end because the shipper holds a new RFP and selects a different primary carrier. Suppose an RFP is held and the primary carrier learns that he has “lost” the RFP, so he will soon lose his primary position. This alters his dynamic incentives. Typically, about four or five weeks pass between the announcement of the RFP outcome and the enactment of the new routing guide that results from that RFP. This means that the carrier experiences a one-month “lame duck” period in which he knows that, after the end of the month, there is no prospect of future relationship surplus. During this lame duck period, we might expect to observe *endgame effects*, where the carrier's tendency to accept loads is diminished. Observing such endgame effects for losing carriers would strongly support the notion that prior to the last few weeks of the contract period,  $\bar{p} > p + \eta$ ; that is, carriers' future relationship surplus induces a cooperative response.

---

<sup>35</sup>See Section 1.2 for a discussion of these three carrier types.

Even for a carrier who “wins” the RFP and will maintain the primary position in the next contract period, dynamic incentives may still be altered in the last few weeks of the current contract period. If a shipper’s incentive scheme were conditioned only on carrier behavior *within a contract period*, then a winning carrier’s dynamic incentives would be greatly lessened by imminent end of the contract period; he would, in effect, get a “free pass,” knowing that the slate will be wiped clean at the start of the next contract period. Observing this kind of endgame effect for winning carriers would therefore not only provide further evidence of a response to dynamic incentives, but would also speak to the *scope of the incentive mechanism* in the time dimension.

While we might worry that the timing of an RFP is not exogenous, we address this concern by restricting our attention to *mass RFP events*, where the shipper holds RFPs simultaneously on at least 30 lanes. We think it is unlikely that poor carrier performance on one lane will affect the shipper’s decision of when to hold an RFP on such a large set of lanes.<sup>36</sup>

To study these hypothesized endgame effects, we estimate a linear probability model

$$\begin{aligned} \text{Accepted}_{sct}^{\ell} = & \beta_0 + \beta_1 (\text{spot rate}_t^{\ell} - \text{contract rate}_{sct}^{\ell}) \\ & + \sum_{k=1}^{18} \alpha_k \mathbb{1}\{k \text{ weeks until end of contract}\} + \epsilon_{sct}^{\ell} \end{aligned} \quad (2)$$

regressing an indicator for the primary carrier’s acceptance of an offer on a set of dummies for the number of weeks until the end of the contract period (when new rates are enacted), along with the deviation profit (the difference between the spot and contract rates), which captures the carriers’ short-run incentives.<sup>37</sup> The pattern of week fixed effects  $\{\alpha_k\}$  over time will provide insight into the proposed end-of-contract effects. As we are interested in the potential endgame effects for both losing (lame duck) carriers and winning carriers, we estimate (2) separately for these two groups of carriers. The estimated coefficients  $\{\hat{\alpha}_k\}$  on the weeks-to-end-of-contract dummies, along with 95% confidence intervals, are plotted in Figure 4.<sup>38</sup>

In the final month of the contract period, losing primary carriers (solid lines) significantly reduce acceptance, with carriers 17pp less likely to accept load offers in the last week (as compared with a baseline rate of 71%). These economically and statistically significant endgame effects provide strong evidence that—prior to learning the RFP outcome—carriers respond to dynamic incentives. The large magnitude of these endgame effects suggests that  $\bar{p} \gg p + \eta$ , which would results from either large relationship rents or harsh shipper rejection penalties.


The estimated coefficients for winning carriers (dashed lines) similarly show a decline (albeit

<sup>36</sup>This approach is intuitively similar to the “mass layoff” approach used to address worker selection issues in the labor literature.

<sup>37</sup>Notation:  $s$  indexes shippers,  $c$  indexes carriers,  $t$  indexes tenders, and  $\ell$  indexes lanes.

<sup>38</sup>The omitted level is  $k = 5$ , i.e. we normalize  $\hat{\alpha}_5 = 0$ .

Figure 4: End-of-contract effects on acceptance probability for winning and losing carriers



Notes: This figure plots the estimated coefficients  $\{\alpha_k\}$  from equation (2) with normalization  $\alpha_5 = 0$ . This choice reflects the fact that the RFP outcome is typically announced about 5 weeks before the end of the contract period. Also labeled in this plot is the approximate timing of the shipper announcing to carriers that an RFP will be held (6-8 weeks prior to the announcement of the RFP outcome).

a smaller one) in acceptance after the RFP outcomes are announced. This happens despite the fact that the winning carriers now know they will continue to be primary carriers. This is consistent with the winning carrier anticipating that the slate will be wiped clean at the start of the new contract period. Once his primary status in the next contract period is secured, the carrier's incentive to perform well in the current contract period is much weakened. This strongly suggests that the scope of the incentive mechanism is *within* rather than *across* contract periods.

Another feature of Figure 6, however, does hint at the possibility of an across-contract period incentive mechanism. We see that, 9-11 weeks before the end of the contract period, the carriers who go on to win the RFP are slightly more likely to accept load offers. It is common for 6-8 weeks to pass between the shipper informing carriers of an upcoming RFP and announcing the winner of the RFP. Thus, by 11 weeks before the end of the contract period, carriers are likely aware that an RFP is imminent. If a carrier believed the RFP outcome to be conditional on his acceptance decisions, this would create an extra dynamic incentive to accept, possibly resulting in a kind of *window-dressing effect*. Indeed, we find evidence that carriers can affect RFP outcomes: while a primary carrier's current-period acceptances have a negligible effect on his probability of winning the next RFP, they have a positive effect on the new contract rate conditional on winning. Specifically, we find that a 10pp reduction in rejection rate increases the probability of winning by

only 0.7pp but, conditional on winning, increases the contract premium by 4 cents per mile (2% of the typical spot rate of \$2 per mile).<sup>39</sup>

From the evidence presented in this subsection, we draw two conclusions: First, the temporal scope of the incentive mechanism is largely within the contract period. Second, the evidence on endgame effects strongly supports the hypothesis that carriers respond to dynamic incentives. This second conclusion, however, comes with a caveat: this evidence on dynamic incentives is limited to a selected subset of relationships (those with mass RFP events) and to a selected time period (the last 18 weeks of the contract period). To further support our conclusion, Section 5.3 will present additional evidence of carriers responding to dynamic incentives based on acceptance decisions across and throughout all relationships in our sample.

## 5.2 Empirical Evidence: Shippers' Strategies

Having established the temporal scope of the incentive mechanism, we now assess its spatial scope, as well as its magnitude. To do this, we study shippers' decisions to demote primary carriers to determine whether (and to what degree) such demotion decisions on lane  $\ell$  are conditioned on the carrier's performance on all lanes within the shipper-carrier relationship or only on the carrier's performance on lane  $\ell$ . Rather than assuming that this spatial scope is the same across all shipper-carrier relationships, we estimate the shipper's demotion strategy separately for large asset-based carriers (large ABCs), small asset-based carriers (small ABCs), and brokers.<sup>40</sup> While we find that the incentive mechanism for brokers is at the firm-to-firm level, we find that the incentive mechanism for large ABCs is at the narrower lane level. For small ABCs, we do not find evidence of punishment.

At a high level, our approach is a simple one: We estimate the following linear probability model of the shipper's demotion strategy:

$$\begin{aligned} \text{Demotion}_{sct}^{\ell} = & \gamma_0 + \gamma_{\text{Rej}(\ell)} \text{Rejection rate}_{sct}^{\ell} + \gamma_{\text{Rej}(-\ell)} \text{Rejection rate}_{sct}^{-\ell} \\ & + \gamma_{\text{Rej}(\ell) \times \text{Rej}(-\ell)} \text{Rejection rate}_{sct}^{\ell} \times \text{Rejection rate}_{sct}^{-\ell} \\ & + \gamma_X X_{sct}^{\ell} + \gamma_{\text{Rej}(\ell) \times X} \text{Rejection rate}_{sct}^{\ell} \times X_{sct}^{\ell} + \epsilon_{sct}^{\ell}, \end{aligned} \quad (3)$$

where  $\text{Demotion}_{sct}^{\ell}$  is an indicator for primary carrier  $c$  being demoted from primary status on lane  $\ell$  between load  $t$  and load  $t + 1$ , and  $X_{sct}^{\ell}$  is a vector that includes the time-invariant rela-

<sup>39</sup>This is a sizable effect, in light of the fact that profit margins are generally low in the trucking industry, at around 2.5% to 6%. See Table 7 in Appendix C for a detailed description of our estimation.

<sup>40</sup>To define carrier types, we use an NMFTA crosswalk to convert the Standard Carrier Alpha Code (SCAC) identifiers in our microdata to US DOT codes. We then match US DOT codes to carriers' DOT registration for the year 2020. This method matches 90% of carriers in our data set to a fleet size variable and a broker/non-broker indicator.

tionship characteristics, along with the spot rate at the time of load  $t$ .<sup>41</sup> The critical regressors are Rejection rate $_{sct}^{\ell}$  and Rejection rate $_{sct}^{-\ell}$ , which capture, respectively, the frequency of rejections by carrier  $c$  on lane  $\ell$  and on all other lanes on which  $c$  is the primary carrier of shipper  $s$ . The coefficients on these variables and their interaction speak to the spatial scope of the incentive mechanism.

**Defining rejection rates** In contrast to the parsimonious model in Section 4, our empirical analysis allows shippers to have memories longer than one load. We adopt a functional form that allows the shipper's strategy to condition on a rejection rate index that summarizes the entire history of rejections in the relationship, though potentially giving greater weight to more recent rejections than less recent ones. For a shipper  $s$ , lane  $\ell$ , and carrier  $c$ , this index takes the following form:

$$\text{Rejection rate}_{sct}^{\ell} = \frac{\sum_{k=0}^{t-1} \alpha^{\text{days}(t-k,t)} \text{Rejection}_{sct-k}^{\ell}}{\sum_{k=0}^{t-1} \alpha^{\text{days}(t-k,t)}}, \quad (4)$$

where  $\text{Rejection}_{sct}^{\ell}$  is an indicator for carrier  $c$  rejecting a load  $t$  from shipper  $s$  on lane  $\ell$ ;  $\text{days}(t-k, t)$  indicates the number of days between load  $t-k$  and load  $t$ ; and  $\alpha \in [0, 1]$  is a daily decay rate.<sup>42</sup> The other-lanes rejection rate  $\text{Rejection}_{sct}^{-\ell}$  is defined analogously using acceptance/rejection decisions by carrier  $c$  on lanes other than lane  $\ell$  on which  $c$  is the primary carrier for  $s$ .

**Relationship characteristics** In estimating the shipper's strategy, we control for relationship characteristics relevant to the payoffs of the shipper and/or carrier.

First and foremost,  $X_{sct}^{\ell}$  includes the *log of average monthly volume*, which is a proxy for  $\delta$ , the frequency of interactions between shipper  $s$  and carrier  $c$  and lane  $\ell$ . We measure monthly volume as the number of loads offered by the shipper on that lane in an active month.<sup>43</sup>

Second,  $X_{sct}^{\ell}$  includes two measures of the *inconsistency of load timing*: The first measures the degree to which the number of weekly loads varies from week to week. The second measures the degree to which the distribution of loads across days of the week varies from week to week. We treat these measures of load consistency as a component of the carrier's match-specific gain ( $\eta_1$ ).<sup>44</sup> Intuitively, if the timing of loads is more consistent, it is easier for the primary carrier to

<sup>41</sup>Recall that by demotion, we mean a change to the lane- $\ell$  routing guide within a contract period that results in carrier  $c$  losing his primary position and being replaced by a new primary carrier. Since this definition is limited to changes within contract periods, any change in primary carrier that coincides with a change in rates (i.e., an RFP) on lane  $\ell$  is not considered a demotion in our analysis.

<sup>42</sup>Note that the special case  $\alpha \downarrow 0$  corresponds to the single-load memory restriction imposed in the model.

<sup>43</sup>By not conditioning on the identity of the accepting carrier, this measure avoids potential endogeneity issues. Moreover, since the primary carrier is the first carrier to receive an offer for each load, this measure approximates the expected number of offers the primary carrier receives per month during the relationship.

<sup>44</sup>See Appendix B for a detailed description of how the two inconsistency measures are constructed.

plan his network of truck movements around the expected timing of offers. As discussed in Section 1.2, in the context of the trucking industry, such network planning could be important for reducing wasteful expenditures on fuel and labor.<sup>45</sup>

Finally,  $X_{sct}^\ell$  includes the difference between the prevailing spot rate for lane  $\ell$  at time  $t$  and carrier  $c$ 's contract rate on lane  $\ell$ . This captures the carrier's short-run incentive for deviation.

**Identification strategy** In estimating equation (3), we face a potential identification challenge stemming from the fact that the shipper's match-specific value ( $\psi$ ) may shape the shipper's optimal strategy but is unobserved and therefore omitted. Several variables on the right-hand side of (3) are endogenous and are likely correlated with  $\psi$ . First, the rejection rate is the result of the carrier's endogenous response to the shipper's strategy. If the shipper's strategy is shaped by the omitted  $\psi$ , then the rejection rate will also be a function of  $\psi$  and therefore correlated with the regression error  $\epsilon_{sct}^\ell$ . Second, the contract rate, which enters  $X_{sct}^\ell$ , is an endogenous outcome of the RFP process and is likely to be positively correlated with the shipper's match-specific value.

To address these endogeneity concerns, we use an instrumental variables approach that exploits exogenous variation in spot rates. First, we instrument for past acceptance/rejection decisions using the spot rates at the time at which each acceptance/rejection decision was made. To that end, we construct an index of past spot rates analogous to the construction of the rejection rate index, which serves as an instrument for the rejection rate.<sup>46</sup> We likewise construct an index of past spot rates on other lanes, which serves as an instrument for the other-lanes rejection rate. Second, we instrument for the contract rate using the spot rate at the time of the RFP in which the contract rate was established.<sup>47</sup> The idea is that at the RFP stage, the current spot rate serves as a competitive pressure on proposed contract rates.<sup>48</sup>

Using past spot rates as instruments is attractive because variation in spot rates is plausibly exogenous for our purposes. While endogenous factors at the market level do determine spot rates, the industry is competitive enough that no one shipper or carrier has significant power to influence them. To satisfy the exclusion restriction, however, past spot rates must not directly affect the shipper's strategy. An identifying assumption is therefore that only the period- $t$  spot rate affects the shipper's period- $t$  demotion decision, something which aligns with industry experts' descriptions of typical demotion decision processes. According to these experts, shippers track carrier performance and make demotion decisions using a scorecard that records various performance aspects,

<sup>45</sup>This network-planning explanation for carriers valuing consistent timing of load offers is widely accepted in the truckload industry.

<sup>46</sup>Specifically, Past spot index $_{sct}^\ell = \sum_{k=0}^{t-1} \alpha^{\text{days}(t-k,t)} \tilde{p}_{t-k}^\ell / \sum_{k=0}^{t-1} \alpha^{\text{days}(t-k,t)}$ .

<sup>47</sup>To be more precise, we use Spot rate $_{sct}^\ell - \text{Spot rate}_{sc0}^\ell$  to instrument for Spot rate $_{sct}^\ell - \text{Contract rate}_{sc}^\ell$ .

<sup>48</sup>Figure 1 shows that the average contract rate tends to adjust with spot rates, though with some lag. This supports the idea that spot rates create competitive pressure on contract rates.



including rejection history, but does not record spot rate history.<sup>49</sup>

**Estimation** We jointly estimate the parameters  $(\alpha, \gamma)$  by GMM. The parameters  $\gamma$  are identified by the standard 2SLS moments. To identify  $\alpha$ , we include a set of additional instruments, the prevailing spot rate on lane  $\ell$  during each of the last five weeks, which, under the exclusion restriction, should be conditionally independent of the outcome.<sup>50</sup> For computational efficiency, we implement this GMM estimation via a nested algorithm. For a given value of  $\alpha$ , the inner step uses 2SLS to obtain estimates of the linear parameters  $\gamma$ , while the outer loop searches for the value of  $\alpha$  that minimizes the GMM objective function.

For the different carrier types, we estimate daily decay rates ranging from  $\hat{\alpha} = 0.9742$  to  $\hat{\alpha} = 0.9992$ . This means that the shipper puts between 2.3% and 54.3% less weight on a rejection one month ago than on a rejection today. Our estimates of the linear parameters  $\hat{\gamma}$  are reported in Table 2. For each carrier type, the parameter estimates for our main specification are the GMM estimates in the second column. For contrast and to illustrate the endogeneity problem described above, we also report in the first column the OLS estimates of the parameters  $\gamma$ .<sup>51</sup> We discuss the GMM estimates and their interpretation below.

**Spatial scope and magnitude of the incentive mechanism** The results in Table 2 show substantial heterogeneity in shipper strategies across carrier types. For brokers, we see that  $\hat{\gamma}_{\text{Rej}(\ell) \times \text{Rej}(-\ell)}$  is positive and statistically significant, indicating that shippers punish multi-lane rejections more harshly. In contrast, for large ABCs, we see that  $\hat{\gamma}_{\text{Rej}(\ell)}$  is positive and significant, while both  $\hat{\gamma}_{\text{Rej}(-\ell)}$  and  $\hat{\gamma}_{\text{Rej}(\ell) \times \text{Rej}(-\ell)}$  are insignificant, indicating that shippers use a single-lane incentive mechanism for this type of carrier. Example 1 suggests a possible explanation for this difference across carrier types. Since an ABC's costs and match-specific benefits on a lane depend in large part on the lane's alignment with the rest of the carrier's network, we would expect large ABCs to have significantly more heterogeneity across lanes than brokers.<sup>52</sup> Such heterogeneity makes

---

<sup>49</sup>Industry experts with whom we discussed these issues include Steve Raetz, Director of Research and Market Intelligence at CH Robinson; other members of Steve's team; and Chris Caplice and Angi Acocella of the MIT Center for Transportation and Logistics.

<sup>50</sup>These five lagged spot rates are instruments for the past acceptance/rejection decisions over the last five weeks. Under the functional form assumption in (4), however, individual rejection decisions enter into the shipper's strategy only through the rejection rate index. Thus, at the true  $\alpha$ , these instruments for the individual acceptance/rejection decisions should be uncorrelated with the error term.

<sup>51</sup>These are the OLS estimates of  $\gamma$  with rejection rate measures constructed using the GMM estimate of the daily decay rate.

<sup>52</sup>The presence or absence of multi-lane punishment might also relate to divides within the shipper's organization. Shippers may divide responsibility for different lanes in various ways, for example, having separate teams managing inbound lanes (shipments of inputs from suppliers) versus outbound lanes (shipments of outputs to customers). Divides like these (combined with a lack of communication between these separate teams) may offer an operational explanation for the shipper's failure to implement multi-lane punishment. Yet such organizational divides arise endogenously; by choosing to separate responsibility for different lanes, shippers forgo the potential benefits of multi-lane punishment.

Table 2: Estimation of shipper's strategy

	All carriers		Large asset-based carriers		Small asset-based carriers		Brokers	
	(OLS)	(GMM)	(OLS)	(GMM)	(OLS)	(GMM)	(OLS)	(GMM)
Rejection rate	0.00842 (0.00048)	0.00233 (0.00467)	0.00823 (0.00082)	0.084 (0.034)	0.012 (0.001)	0.000847 (0.012344)	0.00642 (0.00103)	-0.0151 (0.0110)
Other lane rejection rate	-0.00547 (0.00086)	-0.0166 (0.0111)	-0.00241 (0.00136)	-0.0408 (0.0319)	0.000477 (0.002114)	0.0364 (0.0463)	0.00281 (0.00149)	-0.0754 (0.0230)
Rejection rate × Other lanes rejection rate	0.0188 (0.0014)	0.0743 (0.0216)	0.0137 (0.0022)	0.018 (0.043)	-0.00557 (0.00357)	0.0102 (0.0914)	0.00652 (0.00264)	0.257 (0.055)
Volume	-0.00689 (0.00014)	-0.00684 (0.00058)	-0.00438 (0.00028)	-0.0137 (0.0022)	-0.00415 (0.00036)	-0.000986 (0.002243)	-0.00805 (0.00023)	-0.0106 (0.0006)
Inconsistency (loads / week)	0.0173 (0.0007)	0.0341 (0.00411)	0.0235 (0.0018)	-0.0534 (0.0281)	0.0122 (0.0019)	0.00311 (0.00713)	0.0154 (0.0010)	0.0328 (0.0027)
Inconsistency (day of week)	-0.00984 (0.00044)	-0.0108 (0.0013)	-0.00522 (0.00086)	-0.0154 (0.0049)	-0.00656 (0.00120)	0.00406 (0.00364)	-0.0122 (0.0007)	-0.0207 (0.0018)
Spot rate - contract rate	-0.00224 (0.00037)	-0.00947 (0.00199)	-0.00346 (0.00066)	-0.00264 (0.00587)	-0.00362 (0.00085)	-0.014 (0.009)	-0.00302 (0.00077)	0.00958 (0.00660)
Rejection rate × Volume	-0.00759 (0.00029)	-0.00747 (0.00163)	-0.0105 (0.0005)	0.00927 (0.00451)	-0.00803 (0.00072)	-0.0192 (0.0064)	-0.0089 (0.0006)	0.00788 (0.00223)
Rejection rate × Inconsistency (loads / week)	-0.0141 (0.0011)	-0.0504 (0.0093)	-0.0171 (0.0022)	0.0746 (0.0370)	-0.00792 (0.00313)	0.00527 (0.01088)	-0.0147 (0.0019)	-0.0621 (0.0088)
Rejection rate × Inconsistency (day of week)	-0.00709 (0.00101)	0.00279 (0.00464)	-0.0116 (0.0017)	-0.0103 (0.0129)	-0.0093 (0.0029)	-0.0478 (0.0128)	-0.0064 (0.0020)	0.0257 (0.0081)
Rejection rate × (Spot rate - contract rate)	0.000754 (0.000769)	0.0208 (0.0071)	0.00195 (0.00112)	-0.0287 (0.0191)	0.00347 (0.00199)	0.0252 (0.0249)	0.00358 (0.00187)	-0.0333 (0.0225)
$\hat{\alpha}$	-	1.000	-	0.974	-	0.999	-	0.985
$N$	680,229	680,229	173,787	173,787	67,508	67,508	250,197	250,197

Notes: Standard errors are in parentheses. For ease of interpretation, covariates that are interacted with rejection rate (other lanes rejection rate, volume, inconsistency, and spot rate minus contract rate) are normalized to have mean zero. The GMM specification jointly estimates  $\alpha$  and the linear coefficients presented in this table. An outer loop searches over values of  $\alpha$ , while, for a given  $\alpha$ , an inner step estimates the linear coefficients by 2SLS and computes the 2SLS objective function. For each carrier type, the OLS specification takes the value of  $\alpha$  estimated in the GMM specification as given and estimates the linear coefficients by OLS.

it harder to effectively combine incentives across lanes using simple incentive schemes, such as conditioning relationship continuation across relationships on a joint scorecard.

For all carrier types, the estimated coefficients suggest that the shipper's punishment scheme is soft, rather than harsh. For large ABCs, for instance, our estimate  $\hat{\gamma}_{\text{Rej}(\ell)}$  is positive and very statistically significant, indicating that shippers punish rejections with an increased probability of demotion. At first glance, however, this coefficient may appear very small, as it indicates that an increase in the rejection rate from 0% to 100% increases the probability of demotion between load  $t$  and load  $t + 1$  by only 8.5 percentage points. However, this coefficient should be interpreted in light of the fact that  $\text{Rejection rate}_{sct}^{\ell}$  is a persistent state variable; since  $\hat{\alpha} \gg 0$ , a rejection of one load results in a sustained increase in the probability of demotion for many periods to come.

To get a sense of the economic significance of the estimated degree of punishment, we run a simple simulation to illustrate the effect of a rejection on the expected duration of a relationship.<sup>53</sup>

<sup>53</sup>See Appendix B.2 for the details of our simulation exercise.

The results show that if a large asset-based carrier rejects the first offer for the relationship, the expected relationship duration is 58.5 loads, as compared to 63.5 loads if he accepts the first offer. This 5 load difference is economically large. The mean per-load payment from shipper to carrier is \$1,129, so a rejection early in the relationship may cost the carrier as much as \$5,645 in revenue. The prospect of such a loss from a single rejection is likely to create meaningful incentives for carrier cooperation.

### 5.3 Empirical Evidence: Carriers' Acceptance

To conclude our empirical evidence, we use carrier behavior throughout the relationship to build upon the evidence presented in Section 5.1, bolstering the claim that carriers respond to dynamic incentives and quantifying the magnitude of this response. To do this, we estimate the response of carriers' tendency to accept loads to relationship characteristics. We show that carriers respond strongly to lane volume, which would not be true of a carrier playing static best response.

As we did for shippers' strategies in the previous subsection, we estimate a linear probability model,

$$\begin{aligned} \text{Accepted}_{sct}^{\ell} = & \beta_0 \text{controls}_s^{\ell} + \beta_{\text{volume}} \text{volume}_s^{\ell} \\ & + \beta_{\text{inc}} \text{inconsistency}_s^{\ell} + \beta_{\text{volume} \times \text{inc}} \text{volume}_s^{\ell} \times \text{inconsistency}_s^{\ell} \\ & + \beta_{\text{spot}} (\text{spot} - \text{contract})_t^{\ell} + \beta_{\text{volume} \times \text{spot}} \text{volume}_s^{\ell} \times (\text{spot} - \text{contract})_t^{\ell} + \epsilon_{sct}^{\ell}, \end{aligned} \quad (5)$$

regressing an indicator for carrier  $c$  accepting load  $t$  from shipper  $s$  on lane  $\ell$  on controls, along with a set of lane and relationship characteristics. Our choice of functional form—in particular, the inclusion of interactions between volume and other characteristics—reflects insights from Proposition 1.

**Identification strategy** In estimating (5), we face two identification challenges stemming from the fact that a component of match-specific gain ( $\eta_2$ ), as well as the carrier's cost distribution ( $F$ ), is unobserved and thus omitted.

First, as in the previous subsection, we face the problem that the contract rate is an endogenous object likely correlated with the omitted variable.<sup>54</sup> We again address this issue by instrumenting for the contract rate using the spot rate at the time of the RFP in which the contract rate was established.

Second, while our use of average volume as a proxy for the frequency of future interactions (and thus for the discount factor  $\delta$ ) is in keeping with the empirical literature on relational contracting

<sup>54</sup>Since contract rates are established through an RFP process, a carrier with high match-specific value would tend to submit a lower bid, so we expect upward bias in the OLS estimates of  $\beta_{\text{spot}}$  and  $\beta_{\text{volume} \times \text{spot}}$ .

(e.g., Gil and Marion (2013)), we would face a potential identification challenge if this measure of volume were correlated with the unobserved component of carrier's match-specific gains,  $\eta_2$ . Such correlation might arise either because of a *selection effect* or an *investment effect*. The selection effect would result if carriers with better match-specific value or cost were systematically more likely to be primary carriers on higher-volume lanes.<sup>55</sup> The investment effect would result if the carrier could, for example, adjust his network of truck movements to better serve the contracted lane; he would have greater incentive to carry out such adjustments on a higher-volume lane.<sup>56</sup>

In either case, a positive correlation between volume and the unobserved match-specific value would induce bias in our key parameter of interest,  $\beta_{\text{volume}}$ . To address the potential bias resulting from the selection effect, we include two sets of fixed effects that absorb variation in the carrier's unobserved match-specific gain and cost: the first are shipper-carrier fixed effects; the second are carrier-origin-destination-year fixed effects, where origin and destination are defined as Census regions. By including the shipper-carrier fixed effects, we absorb a variety of potential shipper-carrier specific components of the unobserved value, including, for instance, relationship-specific knowledge and integration of payment or communication systems.<sup>57</sup> By including the carrier-origin-destination-year fixed effects, we absorb key geographic components of the carrier's match-specific value, namely the compatibility of a particular (broadly defined) route with the rest of the carrier's network.

For each carrier type, the first column in Table 3 reports the OLS estimates of Equation (5) while the second column reports the IV estimates. The fifth column reports estimates for the main specification, IV with carrier-origin-destination-year and shipper-carrier fixed effects included. For comparison, the third and fourth columns each report IV estimates for a specification including only one of these sets of fixed effects.

**Response to dynamic incentives** The results for the main specification indicate that  $\beta_{\text{volume}}$  is positive and significant for all carrier types, which would not be true of a carrier playing static best responses.<sup>58</sup> We interpret this result as strong evidence that carriers respond to dynamic incentives. We also observe striking differences in the magnitude of this response across carrier types: large

<sup>55</sup>There are several ways this selection could arise. First, carriers with high  $\eta$  could be more likely to submit bids for RFPs on high-volume lanes. Second, in choosing winners of RFPs, shippers might be more likely to select high- $\eta$  carriers from among the bidders on high-volume lanes.

<sup>56</sup>Note that this investment effect story involves a violation of the timing assumption of the model, which imposes that  $(\eta_2, F)$  are fixed and known at time  $t = 0$ .

<sup>57</sup>Note, however, that when these fixed effects are included, we estimate the parameters by exploiting variation in characteristics *across* different lanes of the same shipper-carrier pair. This approach would not produce meaningful estimates of the effects of these characteristics if shippers employed multilane punishment strategies; if that were the case, a carrier's acceptance/rejection decisions on lane  $\ell$  might respond to characteristics of other lanes. However, our analysis indicates that this is not a concern for ABCs.

<sup>58</sup>For such a carrier, acceptances would depend only on the current spot rate and threshold  $\underline{p} = \eta_1 + \eta_2 + p$ .

Table 3: Estimation of carriers' acceptance

	All carriers					Large asset-based carriers				
	(OLS)	(IV)	(IV-FE1)	(IV-FE2)	(IV-FE3)	(OLS)	(IV)	(IV-FE1)	(IV-FE2)	(IV-FE3)
Volume	0.00792 (0.00057)	0.000376 (0.000763)	0.00790 (0.00167)	0.0188 (0.0009)	0.00925 (0.00157)	-0.00408 (0.00142)	0.0325 (0.0035)	0.150 (0.029)	0.0558 (0.0079)	0.0893 (0.0169)
Spot rate - contract rate	-0.223 (0.001)	-0.261 (0.003)	-0.150 (0.002)	-0.231 (0.002)	-0.150 (0.002)	-0.142 (0.002)	-0.191 (0.006)	-0.126 (0.012)	-0.133 (0.005)	-0.117 (0.008)
Inconsistency (loads / week)	-0.0898 (0.0024)	-0.0869 (0.0026)	-0.0686 (0.0028)	-0.0872 (0.0027)	-0.0686 (0.0028)	-0.264 (0.007)	-0.262 (0.010)	-0.173 (0.015)	-0.177 (0.009)	-0.156 (0.011)
Inconsistency (day of week)	-0.0382 (0.0015)	-0.0370 (0.0016)	-0.0346 (0.0021)	-0.0427 (0.0020)	-0.0282 (0.0019)	-0.0576 (0.0035)	0.00157 (0.00663)	0.0995 (0.0277)	-0.0170 (0.0091)	0.0460 (0.0159)
Volume $\times$ (Spot rate - contract rate)	0.0356 (0.0009)	0.362 (0.019)	0.388 (0.030)	0.362 (0.028)	0.334 (0.028)	0.0352 (0.0023)	-0.772 (0.064)	-1.460 (0.311)	-0.774 (0.111)	-0.902 (0.189)
Volume $\times$ Inconsistency (loads / week)	-0.0330 (0.0022)	-0.0178 (0.0025)	-0.0270 (0.0024)	-0.0219 (0.0024)	-0.0264 (0.0023)	-0.0717 (0.0065)	-0.0924 (0.0087)	-0.0592 (0.0113)	-0.125 (0.010)	-0.0592 (0.0085)
Volume $\times$ Inconsistency (day of week)	-0.0589 (0.0016)	-0.0332 (0.0023)	-0.0432 (0.0021)	-0.0300 (0.0017)	-0.0368 (0.0019)	-0.0791 (0.0038)	-0.109 (0.006)	0.0302 (0.0104)	-0.0561 (0.0048)	0.0104 (0.0058)
< 7 days since promtion	-0.0848 (0.0020)	-0.0919 (0.0022)	-0.0330 (0.0019)	-0.0441 (0.0020)	-0.0296 (0.0018)	-0.143 (0.005)	-0.124 (0.007)	0.0127 (0.0150)	-0.0289 (0.0072)	-0.000743 (0.009930)
<i>Fixed effects</i>										
Carrier $\times$ region $\times$ region $\times$ year			X		X			X		X
Shipper $\times$ carrier				X	X				X	X
N	796346	796346	796167	796346	796167	167095	167095	167065	167095	167065

	Small asset-based carriers					Brokers				
	(OLS)	(IV)	(IV-FE1)	(IV-FE2)	(IV-FE3)	(OLS)	(IV)	(IV-FE1)	(IV-FE2)	(IV-FE3)
Volume	-0.0374 (0.0019)	-0.0352 (0.0019)	0.0223 (0.0031)	-0.0422 (0.0027)	0.0194 (0.0031)	0.0230 (0.0009)	0.0299 (0.0013)	0.0277 (0.0012)	0.0473 (0.0013)	0.0292 (0.0013)
Spot rate - contract rate	-0.323 (0.004)	-0.159 (0.008)	-0.0950 (0.0081)	-0.111 (0.008)	-0.0969 (0.0081)	-0.338 (0.02)	-0.356 (0.004)	-0.248 (0.005)	-0.337 (0.004)	-0.243 (0.005)
Inconsistency (loads / week)	-0.183 (0.012)	-0.135 (0.014)	-0.123 (0.016)	-0.191 (0.015)	-0.123 (0.016)	-0.0339 (0.0030)	-0.0278 (0.0035)	-0.0208 (0.0043)	-0.0483 (0.0041)	-0.0248 (0.0043)
Inconsistency (day of week)	-0.113 (0.006)	-0.0850 (0.0070)	-0.00348 (0.01170)	0.00453 (0.009630)	-0.00132 (0.01190)	0.00305 (0.00218)	-0.0638 (0.0070)	-0.0159 (0.0035)	-0.0293 (0.0041)	-0.0179 (0.0038)
Volume $\times$ (Spot rate - contract rate)	-0.0648 (0.0028)	-0.0743 (0.0129)	0.0254 (0.0162)	0.0641 (0.0139)	0.0293 (0.0161)	0.0771 (0.0017)	0.576 (0.049)	0.248 (0.030)	0.385 (0.030)	0.261 (0.032)
Volume $\times$ Inconsistency (loads / week)	-0.0605 (0.0076)	-0.0166 (0.0086)	-0.0742 (0.0097)	-0.0698 (0.0091)	-0.0726 (0.0097)	-0.0255 (0.0030)	-0.0457 (0.0040)	-0.00803 (0.00357)	-0.0205 (0.0038)	-0.0108 (0.0036)
Volume $\times$ Inconsistency (day of week)	-0.0633 (0.0056)	-0.0257 (0.0066)	-0.0299 (0.0088)	0.00580 (0.00832)	-0.0224 (0.0090)	-0.0304 (0.0024)	-0.0635 (0.0042)	-0.0297 (0.00293)	-0.0199 (0.0029)	-0.0289 (0.0030)
< 7 days since promtion	-0.113 (0.008)	-0.106 (0.008)	-0.0525 (0.0063)	-0.0625 (0.0067)	-0.0544 (0.0063)	-0.0271 (0.0030)	-0.0361 (0.0035)	-0.00188 (0.00266)	-0.00873 (0.00301)	-0.00165 (0.00267)
<i>Fixed effects</i>										
Carrier $\times$ region $\times$ region $\times$ year			X		X			X		X
Shipper $\times$ carrier				X	X				X	X
N	64478	64478	64458	64478	64458	255686	255686	255615	255686	255615

*Notes:* Standard errors in parentheses. Controls include distance and distance squared. Standard errors are in parentheses. For ease of interpretation, the covariates that are interacted with volume (inconsistency and (spot rate - contract rate)) are normalized to have mean zero. Specifications IV-FE1 and IV-FE3 include carrier  $\times$  origin region  $\times$  destination region  $\times$  year fixed effects; specifications IV-FE2 and IV-FE3 include shipper  $\times$  carrier fixed effects. For the former set of fixed effects, regions are determined using US Census regions of loads' origin and destination locations. The small difference in sample size across the specifications with and without fixed effects reflects the fact that singleton observations are dropped in the specifications with fixed effects.

ABCs exhibit a very strong dynamic response (doubling volume increases acceptance probability by 6.2pp), whereas brokers (2.0pp) and small ABCs (1.3pp) show weaker responses. Notably, this ordering of carrier types by responsiveness to dynamic incentives aligns precisely with the relative strength of same-lane dynamic incentives created by the incentive schemes each type faces.<sup>59</sup>

As described above, the appropriateness of our average volume measure as a proxy for the discount factor  $\delta$  is potentially threatened by the selection and investment effects, which, if present, would result in an upward bias in  $\hat{\beta}_{\text{volume}}$ . However, comparing the IV estimates to the IV-FE1 and IV-FE2 specifications, we see that, for ABCs, the inclusion of fixed effects actually *increases*  $\hat{\beta}_{\text{volume}}$ . This speaks against the hypothesized selection and/or investment effects. However, the fact that the inclusion of fixed effects changes the estimates for ABCs so substantially suggests that there is substantial heterogeneity in costs ( $F$ ) and match-specific values ( $\eta_2$ ) which is absorbed by these fixed effects. For brokers, in contrast, the inclusion of fixed effects changes the estimates very little; this suggests that brokers are much less heterogeneous in terms of costs and match-specific values than ABCs, which accords with the fact that brokers are not differentiated by different networks of truck movements.

While we rule out the potential endogeneity issue caused by a positive correlation between volume and relationship-specific investments, these investments—which are likely to take the form of a carrier adjusting his network of truck movements—could still directly affect carriers’ behavior. To assess the role of network adjustments, we include as a regressor an indicator for whether a load  $t$  occurs within 7 days after carrier  $c$ ’s promotion to primary status on lane  $\ell$ . Since network adjustments likely take some time to complete, we might expect the carrier’s tendency to accept to be lower when it is freshly promoted. This is the case only for small ABCs. We find that a small ABC is 5.4pp less likely to accept a load in the first week after promotion.

**Responses to inconsistency and spot rates** In addition to shedding light on responses to dynamic incentives, the estimates in Table 3 also provide insight into two other key aspects of carrier behavior and how these differ by carrier type.

First, we see that inconsistency (a component of  $\eta_1$ ) does indeed affect carriers’ acceptance decisions. In particular, lanes with more inconsistency in the number of loads per week have lower acceptance probabilities, with this effect being stronger on higher-volume lanes (consistent with the predictions of Proposition 1). The measure of inconsistency in the timing of loads within the week is seemingly of lesser importance.

Second, responses to variation in spot rates differ substantially across carrier types, with brokers being far more sensitive than ABCs. This is intuitive, as the profit margins of brokers, who

---

<sup>59</sup>Recall that results in the previous subsection show that large ABCs face a relatively harsh own-lane punishment scheme, while brokers face a multi-lane scheme and small ABCs face a negligible degree of punishment.

subcontract loads to the spot market, are much more closely tied to spot rates than those of ABCs, who transport loads using their own physical assets. Note that since spot rates in reality exhibit autocorrelation, lower acceptance in response to a higher current spot rate captures two effects: (i) contemporaneous temptation and (ii) a decrease in the carrier's continuation value from the relationship (since future spot rates are also likely to be higher).

## 6 Discussion

This section compares the punishment mechanism proposed in Section 4 with two alternative mechanisms: learning (Section 6.1) and learning coupled with investment (Section 6.2). Taking stock of all of our empirical findings (Section 6.3), we conclude that punishment is the main mechanism for brokers and large ABCs, while learning coupled with investment is the main mechanism for small ABCs.

### 6.1 Learning

The fact that rejections increase the likelihood of demotion could be explained by either a punishment mechanism or a learning mechanism. Suppose carriers have unobserved characteristics, such as relationship-specific gains or costs. Then, past rejections would indicate a higher likelihood of future rejections, thereby affecting the shipper's expected value from maintaining the carrier's primary status. Note that the primary (first-ranked) carrier enjoys greater consistency in the timing of offered loads, facilitating network planning and load fulfillment. Thus, the opportunity cost of maintaining a primary carrier who is likely to reject offers is the higher acceptance probability that the first backup carrier *would have* were he to be primary. This is one reason why the shipper might prefer a primary carrier with higher acceptance probability. From the perspective of a primary carrier, the shipper learning from past rejections would create dynamic incentives for the carrier to accept more loads similar to if the shipper were playing a punishment strategy. However, another pattern of behavior distinguishes these two mechanisms.

**Assumption 1.** *Each carrier has a permanent tendency to accept loads and the shipper holds independent priors over carriers' acceptance tendencies.*

Under Assumption 1, the potential learning mechanism described above can be simplified to a bandit problem with independent arms, a common model in the literature on learning. By choosing a carrier to be the primary carrier and receive the first offer in each period, the shipper gradually learns that carrier's tendency to accept loads as a primary carrier. The solution of the shipper's dynamic optimization problem is as follows: each period, she chooses as primary the carrier with

Table 4: Probability of repromotion

	Brokers		Large ABCs		Small ABCs	
	Full sample	First half	Full sample	First half	Full sample	First half
Return	116	61	106	65	38	18
Never again	4125	2035	1378	655	647	180
Return probability	2.74%	2.91%	7.14%	9.03%	5.55%	9.09%

*Notes for Table 4:* This table presents evidence on the frequency of a carrier who is demoted from the primary position on a lane  $\ell$  ever returning to the primary position on lane  $\ell$ . The first row indicates the number of times in the given subsample that we observe such a return. The second row lists the number of instances where the demoted carrier never again regains the primary position on the lane. One concern with this exercise is that—since our sample is finite—we might not observe a carrier regaining the primary position because this occurs after the end of our sample period. To ameliorate this concern, for each carrier type, we consider two subsamples: the “full sample” column includes all demotions that occur on lanes with a single primary carrier at any point in time during the 2015-2019 period, the “first half” column further restricts to demotions that occur in the 2015-2017 period.

the highest Gittins index, which captures both the exploitation and exploration value of choosing a given carrier over the outside option.<sup>60</sup> Given that the tendency to accept loads should be independent across carriers conditional on observed characteristics, a carrier’s Gittins index evolves only when he is chosen as primary.

A key prediction of this learning model is the prevalence of a “switch-back pattern.” Consider, for example, three carriers—Carrier 1, Carrier 2, and Carrier 3—ordered according to their initial Gittins indices. Initially, Carrier 1 is primary. The shipper continues to choose Carrier 1 until her belief about this carrier’s tendency to accept loads drops just below that of Carrier 2, at which point she switches to Carrier 2. While the shipper chooses Carrier 2, the Gittins index of Carrier 1 remains constant and well above that of Carrier 3. Thus, the next time the shipper needs to make a switch, she switches back to Carrier 1 rather than switching to Carrier 3. This intuition generalizes: when learning steps are small, we expect to see a shipper switching from Carrier 1 to Carrier 2 and then back to Carrier 1 much more often than we see her switching from Carrier 1 to Carrier 2 and then to Carrier 3.<sup>61</sup>

While the learning model predicts the prevalence of the switch-back pattern, it is rare for us to observe a demoted carrier later being repromoted in the data. Table 4 breaks down the probabilities of repromotion by carrier type and timing of demotion. For an asset-based carrier who is demoted in 2015-2017, the probability of ever being repromoted is only about 9% (see Table 4).<sup>62</sup> The repromotion probability for brokers is even lower, at less than 3%.<sup>63</sup> The fact that the switch-back

<sup>60</sup>The exploitation value of an option refers to the expected payoff of that option given the current beliefs. The exploration value refers to the informational value of an additional observation of that option. See Whittle (1980) and Weber et al. (1992) for details.

<sup>61</sup>See Appendix B.2.2 for the details of a multi-arm bandit model and an illustration of the switch-back patterns.

<sup>62</sup>Conditional on being repromoted, the average time between demotion and repromotion is about 200 days.

<sup>63</sup>This is consistent with the lack of heterogeneity in brokers’ performance; there is little to learn about individuals drawn from such a homogeneous population.



pattern is so rarely observed for all carrier types serves as strong evidence against learning as the sole mechanism driving the dynamics in shipper-carrier relationships.

## 6.2 Learning coupled with investment

Though our data rejects learning as the only mechanism, a combination of learning and relationship-specific investment could potentially explain the behavior of ABCs.

For an ABC, investment in a relationship on lane  $\ell$  would likely take the form of the carrier adjusting his network to accommodate loads on lane  $\ell$ . However, suppose that, when the carrier is demoted on lane  $\ell$ , he undoes this investment. Note that carriers undoing their investments violates Assumption 1.<sup>64</sup> Returning to our previous example with three carriers, suppose now that, upon being demoted, Carrier 1 undoes his investment. Then, when the shipper later contemplates demoting Carrier 2, her value from switching back to Carrier 1 may be lower than the value of continuing with Carrier 1 at the time of the first switch. This means that the shipper might now be inclined to choose Carrier 3, instead of switching back to Carrier 1.

## 6.3 Wrap up

We now take stock of our empirical findings and draw conclusions about the most plausible mechanisms for different carrier types. We revisit six empirical features: (i) permanence of demotion, (ii) spatial scope of dynamic incentives, (iii) heterogeneity in carriers' performance across relationships, (iv) shipper's commitment to demotions, (v) negative endgame effects for winning carriers, and (vi) adjustments in the first seven days since promotion. Table 5 summarizes our empirical findings for these features, as well as which mechanisms—punishment, learning, and learning coupled with investment—each feature is consistent with. As discussed in Section 6.1, the permanence of demotion rules out learning as the sole mechanism underlying shipper-carrier relationships for all types of carriers. Thus, what remains is to compare the punishment mechanism proposed in Section 4 and the learning-and-investment mechanism proposed in Section 6.2.

**Brokers** Our empirical findings for brokers strongly favor a punishment mechanism. In particular, the fact that brokers' dynamic incentives operate at the shipper-carrier level supports a multi-lane punishment mechanism. Moreover, the lack of heterogeneity across brokers' performance rules out any mechanism involving learning, since heterogeneity is a precondition for learning.

---

<sup>64</sup>On the other hand, learning about *permanent* investments, even when these investments take place only after promotion, could be accommodated by the learning model in Section 6.1.

Table 5: Mechanisms and evidence: A summary

	Section	Mechanisms			Empirical evidence		
		Punishment	Learning	Learning + Investment	Brokers	Large ABCs	Small ABCs
Permanent demotion	6.1	✓ or ✗	✗	✓ or ✗	✓	✓	✓
Spatial scope*	5.2	<i>scl</i> or <i>sc</i>	<i>scl</i> or none	<i>scl</i> or none	<i>sc</i>	<i>scl</i>	none
Heterogeneity	5.3 <sup>c</sup>	✓ or ✗	✓	✓	✗	✓	✓
Commitment	6.3 <sup>c</sup>	✓	✗	✗	—	✓	—
Endgame effects (winners)	5.1 <sup>c</sup>	✓ or ✗	✓ or ✗	✗	✓	✓	✗
Investment (network)	5.3	✓ or ✗	✗	✓	—	—	✓

*Notes:* The Section column indicates the section of the paper which presents empirical evidence on each feature. A superscript *c* indicates that additional evidence on this feature is presented in Appendix C.

\* Spatial scope here refers to the spatial scope of the linkage between acceptance and demotion, i.e., whether a shipper demotes a carrier on lane  $\ell$  because of rejections on lane  $\ell$  or also because of rejections on other lanes.

*Key:* *scl* indicates lane-level relationship scope; *sc* indicates firm-to-firm (i.e., shipper-carrier) level relationship scope; none indicates no evidence of demotion resulting from rejections. In the Mechanisms columns, ✓ indicates that the mechanism is consistent with the feature, whereas ✗ indicates that the mechanism is not consistent with the feature. ✓/✗ indicates that the mechanism is consistent with both the presence and the absence of the feature. In the Evidence columns, ✓ indicates that we observe the feature, and ✗ indicates that we observe the absence of the feature. — indicates that there is no evidence either way.

**Large ABCs** While our evidence on spatial scope and heterogeneity for large ABCs carriers is consistent with both a punishment mechanism and a learning-and-investment mechanism, other evidence is consistent with only the former. In particular, we find strong evidence that, in their relationships with large ABCs, shippers commit to a demotion strategy. To test for shippers' commitment, we compare the acceptance rates and contract rates of three groups of carriers: (a) RFP winners who are never demoted, (b) RFP winners who are eventually demoted, and (c) non-RFP winners who are promoted. For outcomes  $y \in \{\text{Acceptance, Contract rate}\}$  we estimate

$$y_{sct}^{\ell} = \sum_{g \in \mathcal{G}} \beta_g \mathbb{1}\{g_{sct}^{\ell} = g\} + \underbrace{\gamma (\text{Spot rate}_t^{\ell} - \text{Contract rate}_{sct}^{\ell})}_{\text{only included for } y=\text{Acceptance}} + \epsilon_{sct}^{\ell},$$

where  $\mathcal{G} = \{(a), (b), (c)\}$  is the set of the three groups of primary carriers described above. Evidence that on average, group (b) is better than group (c) in both acceptance rates and contract rates would indicate that a shipper is, on average, made worse off by her decision to demote the primary carrier; this would be strong evidence against any mechanism that involves learning. Indeed, for large ABCs, we find that promoted carriers have substantially worse acceptance rates (15 pp lower)

than demoted carriers without any accompanying discount in contract rate.<sup>65</sup> This, along with the lack of evidence of investment, supports our conclusion that punishment is the main mechanism for large ABCs.

**Small ABCs** In contrast to brokers and large ABCs, small ABCs exhibit behavior at the beginning and end of relationships that suggests a role for relationship-specific investment.

First, Section 5.3 provides evidence on carriers’ performance at the beginning of relationships. Regression (5) includes an indicator for whether a carrier was promoted to primary status on the lane less than seven days ago. The coefficient on this regressor is significant only for small ABCs, who accept about 5pp less often in this window. This “slow start” in the relationship is consistent with small ABCs undertaking investments that take time to be fully realized.

Second, our evidence on negative endgame effects in Section 5.1 is also relevant. If relationship-specific investments were the key driver for carriers’ acceptance, we would not expect winning carriers’ performance to drop significantly after the announcement of the auction outcome for the next contract period. Since winning carriers will continue to service the contracted lane, they should have no reason to undo their investments. Our endgame effects analysis broken down by carrier type shows evidence of decreasing performance—and therefore evidence against the investment story—only for brokers and large ABCs.<sup>66</sup>

The potential importance of investment for small ABCs’ performance leaves open the possibility that a learning-and-investment mechanism drives the dynamic interactions between shippers and small ABCs. Since we do not observe carriers’ networks and therefore cannot provide direct evidence of network adjustment, this paper does not explore this potential mechanism in detail; this would, however, be an interesting avenue for future research on small ABCs.

## 7 Conclusion

In this paper, we ask how informal interfirm relationships work in an economically important setting: the US truckload freight industry. We use a novel transaction-level data set uniquely well-suited to studying informal relationships to provide evidence on the mechanism governing these relationships, as well as the scope of that mechanism.

We begin by presenting evidence of endgame effects, a phenomenon that suggests that the temporal scope of the incentive mechanism is within contract periods. Next, we estimate the shipper’s demotion strategy. The results indicate that while the spatial scope of the demotion strategy is

---

<sup>65</sup>Table 8 in Appendix C compares the acceptance rates and contract rates for each group of carriers. The result is inconclusive about the commitment of brokers and small ABCs. For brokers, promoted carriers tend to have higher acceptance rate. For small ABCs, promoted carriers tend to have lower contract rates.

<sup>66</sup>See Figure 7c in Appendix C.

limited to a single lane for large asset-based carriers, shippers employ multi-lane punishment for brokers. Third, we quantify carriers' responses to the dynamic incentives generated by the incentive scheme. Finally, we address alternative, non-punishment mechanisms. Our evidence suggests that, while punishment seems to be the primary mechanism at play, other mechanisms—in particular, specific investment—likely play a role for relationships involving small asset-based carriers.

Taken together, these results provide valuable insight for future empirical work both on long-term relationships and on the trucking industry. First, we use rich microdata to empirically test the common assumption that relationship scope is at the firm-to-firm level. Our findings illustrate the value of such a test, which we demonstrate to be feasible with increasingly available microdata. Second, by providing a detailed empirical description of the shipper-carrier relationships around which the trucking industry is organized, this paper serves as a key stepping stone to studying other important questions about this macroeconomically vital, yet understudied, industry.

## References

- Acocella, A. and C. Caplice (2023). Research on truckload transportation procurement: A review, framework, and future research agenda. *Journal of Business Logistics*.
- Acocella, A., C. Caplice, and Y. Sheffi (2020). Elephants or goldfish?: An empirical analysis of carrier reciprocity in dynamic freight markets. *Transportation Research Part E: Logistics and Transportation Review* 142, 102073.
- Andrews, I. and D. Barron (2016). The allocation of future business: Dynamic relational contracts with multiple agents. *American Economic Review* 106(9), 2742–59.
- Baker, G., R. Gibbons, and K. J. Murphy (2002). Relational contracts and the theory of the firm. *The Quarterly Journal of Economics* 117(1), 39–84.
- Baker, G. P. and T. N. Hubbard (2003). Make versus buy in trucking: Asset ownership, job design, and information. *American Economic Review* 93(3), 551–572.
- Baker, G. P. and T. N. Hubbard (2004). Contractibility and asset ownership: On-board computers and governance in us trucking. *The Quarterly Journal of Economics* 119(4), 1443–1479.
- Barron, D., R. Gibbons, R. Gil, and K. J. Murphy (2020). Relational adaptation under reel authority. *Management Science* 66(5), 1868–1889.
- Bernheim, B. D. and M. D. Whinston (1990). Multimarket contact and collusive behavior. *The RAND Journal of Economics*, 1–26.

- Blouin, A. and R. Macchiavello (2019). Strategic default in the international coffee market. *The Quarterly Journal of Economics* 134(2), 895–951.
- Board, S. (2011). Relational contracts and the value of loyalty. *American Economic Review* 101(7), 3349–67.
- Brugues, F. (2023). Take the goods and run: Contracting frictions and market power in supply chains. Technical report, Working Paper.
- Caplice, C. (2007). Electronic markets for truckload transportation. *Production and Operations Management* 16(4), 423–436.
- Gil, R., M. Kim, and G. Zanarone (2021). Relationships under stress: Relational outsourcing in the us airline industry after the 2008 financial crisis. *Management Science*.
- Gil, R. and J. Marion (2013). Self-enforcing agreements and relational contracting: evidence from california highway procurement. *The Journal of Law, Economics, & Organization* 29(2), 239–277.
- Green, E. J. and R. H. Porter (1984). Noncooperative collusion under imperfect price information. *Econometrica: Journal of the Econometric Society*, 87–100.
- Harris, A. and T. M. A. Nguyen (2023). Long-term relationships and the spot market: Evidence from us trucking.
- Hortaçsu, A., O. R. Natan, H. Parsley, T. Schwieg, and K. R. Williams (2021). Organizational structure and pricing: Evidence from a large us airline. Technical report, National Bureau of Economic Research.
- Hubbard, T. N. (2001). Contractual form and market thickness in trucking. *RAND Journal of Economics*, 369–386.
- Levin, J. (2003). Relational incentive contracts. *American Economic Review* 93(3), 835–857.
- Li, J. and N. Matouschek (2013). Managing conflicts in relational contracts. *American Economic Review* 103(6), 2328–51.
- Macaulay, S. (1963). Non-contractual relations in business: A preliminary study. *American sociological review*, 55–67.
- Macchiavello, R. and A. Morjaria (2015). The value of relationships: evidence from a supply shock to kenyan rose exports. *American Economic Review* 105(9), 2911–45.

- Macchiavello, R. and A. Morjaria (2023). Relational contracts: Recent empirical advancements and open questions. Technical report, National Bureau of Economic Research.
- MacLeod, W. B. and J. M. Malcomson (1989). Implicit contracts, incentive compatibility, and involuntary unemployment. *Econometrica: Journal of the Econometric Society*, 447–480.
- Malcomson, J. M. (2010). Relational incentive contracts. *Handbook of Organizational Economics*.
- Masten, S. E. (2009). Long-term contracts and short-term commitment: Price determination for heterogeneous freight transactions. *American Law and Economics Review* 11(1), 79–111.
- Pavan, A., I. Segal, and J. Toikka (2014). Dynamic mechanism design: A myersonian approach. *Econometrica* 82(2), 601–653.
- Rose, N. L. (1985). The incidence of regulatory rents in the motor carrier industry. *The RAND Journal of Economics*, 299–318.
- Rose, N. L. (1987). Labor rent sharing and regulation: Evidence from the trucking industry. *Journal of Political Economy* 95(6), 1146–1178.
- Scott, A., C. Parker, and C. W. Craighead (2017). Service refusals in supply chains: Drivers and deterrents of freight rejection. *Transportation Science* 51(4), 1086–1101.
- Shapiro, C. and J. E. Stiglitz (1984). Equilibrium unemployment as a worker discipline device. *The American economic review* 74(3), 433–444.
- Taylor, D. (2011). Trucking industry updates (03/22/2011).
- Weber, R. et al. (1992). On the gittins index for multiarmed bandits. *The Annals of Applied Probability* 2(4), 1024–1033.
- Whittle, P. (1980). Multi-armed bandits and the gittins index. *Journal of the Royal Statistical Society: Series B (Methodological)* 42(2), 143–149.

## A Proofs

### A.1 Carriers' acceptance

Denote by  $V$  the average discounted expected utility of the carrier from the relationship. Denote by  $V(A)$  and  $V(R)$  the average discounted expected utilities of the carrier at the beginning of period

$t$  following  $d_{t-1} = A$  and  $d_{t-1} = R$ , respectively. We have

$$\begin{aligned} V &= \mathbf{E}_{\tilde{p}_t, c_t} [\max \{ (1 - \delta)(\eta + p - c_t) + \delta V(A), (1 - \delta)(\tilde{p}_t - c_t) + \delta V(R), \delta V(R) \}] \\ &= \delta V(R) + (1 - \delta) \mathbf{E}_{\tilde{p}_t, c_t} \left[ \max \left\{ \eta + p - c_t + \frac{\delta}{1 - \delta} (V(A) - V(R)), \tilde{p}_t - c_t, 0 \right\} \right], \end{aligned} \quad (6)$$

where

$$V(A) = \sigma_0(A)V + (1 - \sigma_0(A))\underline{V}, \quad (7)$$

$$V(R) = \sigma_0(R)V + (1 - \sigma_0(R))\underline{V}. \quad (8)$$

Let  $\bar{p} = \eta + p + \frac{\delta}{1 - \delta} (V(A) - V(R))$  and  $h(\bar{p}) = \mathbf{E}_{\tilde{p}_t, c_t} [\max \{ \bar{p} - c_t, \tilde{p}_t - c_t, 0 \}]$ . When  $\tilde{p}_t < \bar{p}$ , the carrier's optimal strategy is to accept whenever  $c_t < \bar{p}$ . When  $\tilde{p}_t > \bar{p}$ , the carrier optimally rejects regardless of the cost draw. Thus, the probability of acceptance at each level of spot rate is  $\Pr(d_t = A | \tilde{p}_t) = \mathbf{1}(\tilde{p}_t < \bar{p})F(\bar{p})$ , increasing in the acceptance threshold  $\bar{p}$ . Manipulating Equations (6), (7) and (8) yields the following fixed point equation of  $\bar{p}$ ,

$$\frac{1 - \delta\sigma_0(R)}{\sigma_0(A) - \sigma_0(R)} (\bar{p} - \eta - p) = \delta(h(\bar{p}) - \underline{V}). \quad (9)$$

**Lemma 1.**  $h' \in [0, 1]$  and  $h'' \geq 0$ .

*Proof.* By the independence of spot rates and cost draws,

$$h(\bar{p}) = G(\bar{p}) \int_0^{\bar{p}} (\bar{p} - c) f(c) dc + \int_{\bar{p}}^{\infty} \int_0^{\tilde{p}} (\tilde{p} - c) f(c) g(\tilde{p}) dcd\tilde{p}.$$

Thus  $h'(\bar{p}) = G(\bar{p})F(\bar{p}) \in [0, 1]$  and  $h''(\bar{p}) = g(\bar{p})F(\bar{p}) + G(\bar{p})f(\bar{p}) \geq 0$ .  $\square$

**Proof of Proposition 1** Notice that if  $\eta + p \in \text{Supp}(G)$ , there is an option value to having contracted offers. That is,  $V > \underline{V}$ . This means that when  $\sigma_0(A) > \sigma_0(R)$ , we have  $V(A) > V(R)$ , and thus  $\bar{p} > \eta + p$ . This completes the proof of part (i).

Next, we exploit Equation (9) to generate predictions on how relationship characteristics and the reward-punishment scheme affect the likelihood of the carrier accepting, as captured by threshold  $\bar{p}$ . Referring to the left-hand-side and the right-hand-side of Equation (9) as *LHS* and *RHS*, we have

$$\frac{\partial(LHS - RHS)}{\partial \bar{p}} = \frac{1 - \delta\sigma_0(R)}{\sigma_0(A) - \sigma_0(R)} - \delta h'(\bar{p}) > 1 - h'(\bar{p}) \geq 0,$$

where the last inequality follows from Lemma 1. Note that since  $\frac{\partial(LHS - RHS)}{\partial \bar{p}} > 0$  for all  $\bar{p}$ ,

Equation (9) has a unique solution  $\bar{p}$ . Also,

$$\frac{\partial(LHS - RHS)}{\partial \delta} = \frac{-\sigma_0(R)}{\sigma_0(A) - \sigma_0(R)}(\bar{p} - \eta - p) - (h(\bar{p}) - \underline{V}) < 0.$$

and

$$\frac{\partial(LHS - RHS)}{\partial \eta} = -\frac{1 - \delta\sigma_0(R)}{\sigma_0(A) - \sigma_0(R)} < 0.$$

Thus it follows from the implicit function theorem that  $\frac{\partial \bar{p}}{\partial \delta} > 0$  and

$$\frac{\partial \bar{p}}{\partial \eta} = \left[ 1 - \frac{\delta(\sigma_0(A) - \sigma_0(R))h'(\bar{p})}{1 - \delta\sigma_0(R)} \right]^{-1} \geq 1.$$

Furthermore, notice that  $\frac{\partial h'(\bar{p})}{\partial \delta} = h''(\bar{p})\frac{\partial \bar{p}}{\partial \delta} \geq 0$ . It follows that  $\frac{\partial^2 \bar{p}}{\partial \delta \partial \eta} \geq 0$ . This completes the proof of part (ii).

Finally, we prove part (iii) of Proposition 1. Rewrite Equation (9) as follows

$$(1 - \delta\sigma_0(R))(\bar{p} - \eta - p) - \delta(\sigma_0(A) - \sigma_0(R))(h(\bar{p}) - \underline{V}) = 0.$$

Applying the implicit function theorem to the above equation yields

$$\frac{\partial \bar{p}}{\partial \sigma_0(A)} = -\frac{-\delta(h(\bar{p}) - \underline{V})}{1 - \delta\sigma_0(R) - \delta(\sigma_0(A) - \sigma_0(R))h'(\bar{p})} > 0,$$

and

$$\frac{\partial \bar{p}}{\partial \sigma_0(R)} = -\frac{-\delta(\bar{p} - \eta - p - h(\bar{p}) + \underline{V})}{1 - \delta\sigma_0(R) - \delta(\sigma_0(A) - \sigma_0(R))h'(\bar{p})} < 0.$$

The first inequality follows because  $h'(\bar{p}) \leq 1$ . For the second inequality, notice in addition that from Equation (9),  $\frac{\bar{p} - \eta - p}{h(\bar{p}) - \underline{V}} = \frac{\delta(\sigma_0(A) - \sigma_0(R))}{1 - \delta\sigma_0(R)} < 1$ .  $\square$

## A.2 Shipper's strategies

First, we derive the shipper's per-period payoff. Each period in a maintained relationship has three possible outcomes: either the carrier accepts the offered load, the carrier rejects because of a high cost draw, or the carrier rejects because of a high spot rate. Thus, the per-period expected utility of the shipper in the relationship equals

$$\begin{aligned} u &= \underbrace{G(\bar{p})F(\bar{p})(\psi - p)}_{\text{accepted}} + \underbrace{G(\bar{p})[1 - F(\bar{p})](-\mathbf{E}[\tilde{p}_t | \tilde{p}_t \leq \bar{p}])}_{\text{rejected because of high cost}} + \underbrace{[1 - G(\bar{p})](-\mathbf{E}[\tilde{p}_t | \tilde{p}_t > \bar{p}])}_{\text{rejected because of high spot rate}} \\ &= -\mathbf{E}[\tilde{p}_t] + G(\bar{p})F(\bar{p})(\psi - p + \mathbf{E}[\tilde{p}_t | \tilde{p}_t \leq \bar{p}]). \end{aligned} \quad (10)$$



Notice that the term  $\mathbf{E}[\tilde{p}_t | \tilde{p}_t \leq \bar{p}] \leq \mathbf{E}[\tilde{p}_t]$  is the shipper's expected payment were she to be served by the spot market conditional on the carrier being willing to accept the offered load. This term represents a selection effect: the carrier has the largest temptation to reject exactly when his acceptance is most valuable to the shipper.<sup>67</sup> This means that even when  $\psi - p > -\mathbf{E}[\tilde{p}_t]$ , a relationship that cannot induce sufficiently high level of cooperation may not be worth sustaining for the shipper. The following lemma provides a sufficient condition for the relationship to be worth sustaining for any incentive scheme with  $0 \leq \sigma_0(R) < \sigma_0(A) \leq 1$ .

**Lemma 2.** *If  $\psi - p + \mathbf{E}[\tilde{p}_t | \tilde{p}_t \leq p] \geq 0$ , then  $u \geq -\mathbf{E}[\tilde{p}_t]$ , that is, the shipper is better off offering to the carrier first than going directly to the spot market.*

*Proof.* Recall that

$$u = -\mathbf{E}[\tilde{p}_t] + G(\bar{p})F(\bar{p})(\psi - p + \mathbf{E}[\tilde{p}_t | \tilde{p}_t \leq \bar{p}]),$$

where  $\mathbf{E}[\tilde{p}_t | \tilde{p}_t \leq \bar{p}]$  represents a selection effect. Define

$$\hat{p} = \inf\{p' \in \text{supp } G : \psi - p + \mathbf{E}[\tilde{p}_t | \tilde{p}_t \leq p'] \geq 0\}.$$

Then  $u > \underline{U}$  and  $\frac{\partial u}{\partial \bar{p}} > 0$  for all  $\bar{p} > \hat{p}$ . The shipper should opt out of the relationship if and only if the sustained level of cooperation satisfies that  $\bar{p} < \hat{p}$ . Under Condition 1 that  $\psi - p + \mathbf{E}[\tilde{p}_t | \tilde{p}_t \leq p] \geq 0$ , we have  $\hat{p} \leq p \leq \bar{p}$ , so the relationship is worth sustaining.  $\square$

We now derive the shipper's average discounted expected utility. Let

$$q = G(\bar{p})F(\bar{p})\sigma_0(A) + (1 - G(\bar{p})F(\bar{p}))\sigma_0(R) \quad (11)$$

denote the probability of maintaining the relationship next period calculated at the beginning of the current period's stage game. The average discounted expected utility  $U$  of the shipper in a maintained relationship is

$$U = (1 - \delta)u + \delta(qU + (1 - q)\underline{U}), \quad (12)$$

where  $\underline{U} = \mathbf{E}[-\tilde{p}_t]$  is the expected payoff of the shipper going directly to the spot market. Thus,

$$U - \underline{U} = \frac{(1 - \delta)(u - \underline{U})}{1 - \delta q}. \quad (13)$$

---

<sup>67</sup>Note that acceptances require both a low spot rate and a low cost draw. However, under the assumption that cost draws are independent of spot rates, as in our model, the hypothetical expected payment in the spot market of an accepted load does not depend on the cost draw being low.

For  $x \in \{\sigma_0(A), \sigma_0(R)\}$ ,

$$\frac{dU}{dx} = \underbrace{\left( \frac{\partial U}{\partial u} \frac{\partial u}{\partial \bar{p}} + \frac{\partial U}{\partial q} \frac{\partial q}{\partial \bar{p}} \right) \frac{\partial \bar{p}}{\partial x}}_{\text{incentive-inducing effect}} + \underbrace{\frac{\partial U}{\partial q} \frac{\partial q}{\partial x}}_{\text{regime-switching effect}}, \quad (14)$$

where  $\partial U/\partial u$ ,  $\partial U/\partial q$ ,  $\partial u/\partial \bar{p}$  and  $\partial q/\partial \bar{p}$  are all positive.

### A.2.1 Proof of Proposition 2

From Equation (14), it follows from  $\partial \bar{p}/\partial \sigma_0(A) \geq 0$  and  $\partial q/\partial \sigma_0(A) \geq 0$  that  $\partial U/\partial \sigma_0(A) \geq 0$ . Thus,  $\sigma_0^*(A) = 1$ . To see that the optimal punishment could be soft, refer to the single-lane relationships in Example 1.

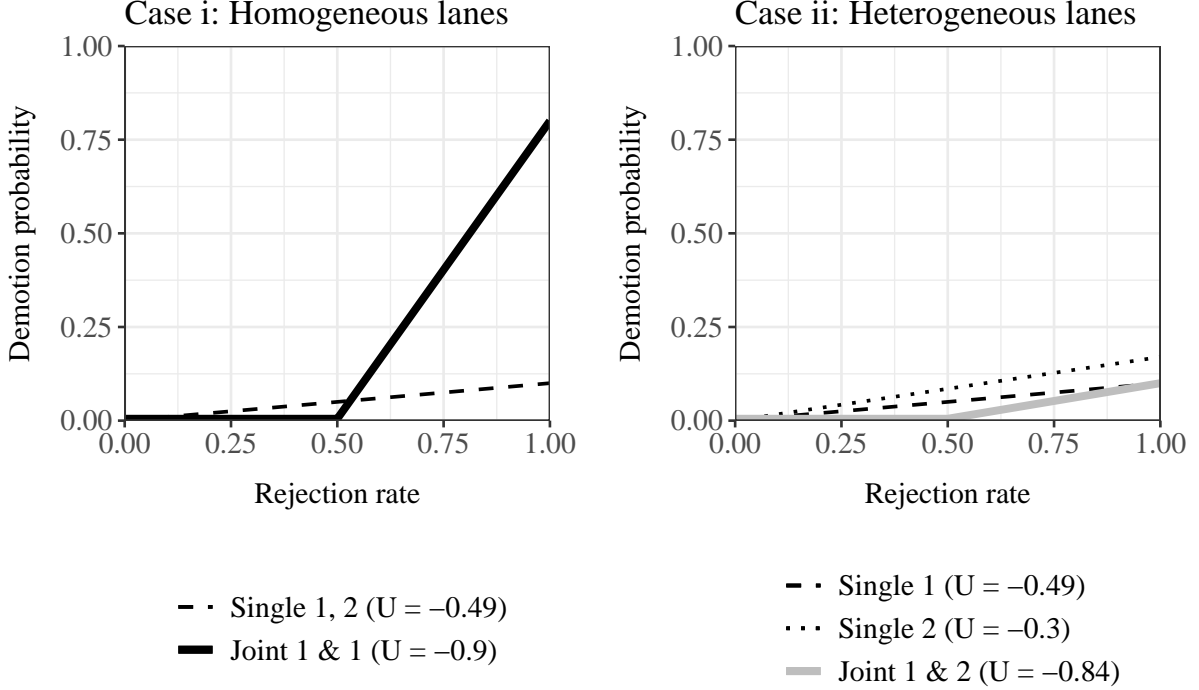
### A.2.2 Details of Example 1

Under the chosen parameter values, the optimal single-lane punishment on lane 1 and lane 2 are soft. Specifically, when  $\eta_1^\ell + \eta_2^\ell + p^\ell = 0.65$ , the optimal demotion probability following a rejection is 0.1, and when  $\eta_1^\ell + \eta_2^\ell + p^\ell = 0.8$ , the optimal demotion probability following a rejection is 0.13. Define a multi-lane rejection rate as the average rejections across all lanes last period,  $\frac{1}{2} \sum_{\ell=1}^2 \mathbf{1}\{d_{t-1} = R\}$ , and consider the shipper's incentive scheme that maps last-period average rejections to a probability of maintaining (or ending) both relationships.

The left panel of Figure 5 plots the optimal single-lane and multi-lane incentive schemes for case i. In this case, the optimal multi-lane incentive scheme forgives non-concurrent rejections but punishes concurrent rejections harshly. Specifically, if the carrier rejects on both lanes, the probability of demotion in the next period is 0.8. This multi-lane incentive scheme benefits the shipper over the optimal single-lane incentive scheme,  $U = -0.9 > 2(-0.49)$ , for two reasons. First, by forgiving non-concurrent rejections, the shipper allows the carrier to attain more allocative efficiency across the two lanes. This, in turn, increases the continuation value of the relationship for the carrier. Second, combining incentives with harsher punishments on joint rejections make these joint rejections low-probability events, and thus, clear signals of noncooperation.

The right panel of Figure 5 plots the optimal single-lane and multi-lane incentive schemes for case ii. The optimal multi-lane incentive scheme in this case takes a similar convex shape as that in case i, but punishes concurrent rejections softly. More importantly, the shipper is strictly worse off by this scheme as compared to using the single-lane optimal incentive scheme,  $U = -0.84 < (-0.49) + (-0.3)$ . This example shows that using a simple incentive scheme that conditions on average rejections (i.e., a common scorecard) across heterogeneous lanes might hurt the shipper.

Figure 5: Gains from pooling homogeneous lanes and losses from pooling heterogeneous lanes



## B Additional empirical details

### B.1 Construction of empirical variables

In this appendix, we explain the construction of right-hand side variables used in our analysis but the details of whose construction is omitted from the main text.

**Inconsistency (loads / week)** Our first empirical measure of the inconsistency of load timing captures a notion of how the number of loads varies from week to week within a month.

Let  $n_{mw}^\ell$  denote the number of loads on lane  $\ell$  in week  $w$  of month  $m$ . We then construct a measure of how much this count varies from week to week within a month by computing  $CV_m^\ell$ , the coefficient of variation among  $(n_{m1}^\ell, n_{m2}^\ell, n_{m3}^\ell, n_{m4}^\ell)$ . We then get a lane-level measure by averaging  $CV_m^\ell$  over all active months on the lane:

$$\text{Inconsistency (loads / week)}^\ell = \frac{1}{M} \sum_m CV_m^\ell.$$

**Inconsistency (day of week)** Our second empirical measure of the inconsistency of load timing captures a notion of how the timing of loads within a week varies from week to week (within a month). For instance, if a lane has 50% of its weekly loads on Monday and 50% of its weekly

loads on Wednesday every week, the lane's loads would be perfectly consistent according to this measure. If, on the other hand, a lane's weekly loads were randomly allocated across days within each week, this lane's loads would be highly inconsistent according to this measure.

Our construction of this measure is motivated by a Chi-squared Goodness of Fit test of the null hypothesis that the lane's distribution of loads across days of the week is the same every week.

For any week  $w$  of month  $m$  and any day of the week  $d$ , the observed fraction of weekly volume on day of the week  $d$  is

$$O_{mwd}^\ell = \frac{n_{mwd}^\ell}{\sum_{d'} n_{mwd'}^\ell}.$$

Across all weeks within month  $m$ , the fraction of volume on day of the week  $d$  is

$$E_{md}^\ell = \frac{\sum_{w'} n_{mw'd}^\ell}{\sum_{w'} \sum_{d'} n_{mw'd'}^\ell}.$$

Under the null hypothesis,  $O_{mwd}^\ell = E_d^\ell$  for all weeks  $w$  and days of the week  $d$ . The Chi-square statistic measures the deviation from this null:

$$\chi_w^\ell = \sum_d \frac{(O_{mwd}^\ell - E_{md}^\ell)^2}{E_{md}^\ell}.$$

Then, we get our lane-level inconsistency measure by averaging across all active months and weeks:

$$\text{Inconsistency (day of week)}^\ell = \frac{1}{M} \sum_m \frac{1}{W} \sum_w \chi_w^\ell.$$

## B.2 Simulation exercises

### B.2.1 Effects of the initial decision on relationship duration

Using the estimated shipper's strategy and the mean acceptance probability, we simulate relationships under two scenarios. In the first scenario, the carrier accepts the first load of the relationship; in the second, he rejects the first load of the relationship. Each load  $t > 1$  is accepted with probability 0.71, the average primary carrier acceptance rate. Based on each acceptance/rejection, we update the rejection rate and compute the probability of demotion between each load  $t$  and  $t + 1$ . At each  $t$ , the relationship continues if the shipper is not demoted (which happens with probability  $1 - \sigma_0(R_t)$ , according to the estimated demotion strategy and the current rejection index state) and if an RFP does not occur. An RFP happens exogenously with constant probability  $\frac{1}{83+1}$ , a frequency that is chosen to match the average number of loads (83) between RFPs observed in the

data. For each scenario, we run 10 million simulations.

### B.2.2 Switch-back patterns in a learning model

Consider a learning model in which a shipper interacts with three carriers: Carrier 1, Carrier 2, and Carrier 3. Each Carrier  $i$  for  $i = 1, 2, 3$  has a permanent tendency to accept loads as a primary carrier,  $q_i \in [0, 1]$ . The shipper holds the independent prior over these carriers' acceptance tendencies, with

$$q_0^1 \sim \text{Beta}(a^1, b^1), \quad q_0^2 \sim \text{Beta}(a^2, b^2), \quad \text{and} \quad q_0^3 \sim \text{Beta}(a^3, b^3),$$

where  $a^i, b^i \in \mathbb{N}^+$  for all  $i = 1, 2, 3$ . In each period, the shipper chooses a carrier as the primary carrier and makes an offer to that carrier. The shipper gets an instantaneous payoff of 1 if the offer is accepted, and 0 if the offer is rejected. Essentially, the shipper faces a multi-armed bandit problem with independent arms, where each arm  $i$  is choosing Carrier  $i$  as primary.

**Evolution of beliefs** Let  $n_t^i$  denote the number of times Carrier  $i$  is chosen and  $k_t^i$  the number of times Carrier  $i$  accepts by time  $t$ . The shipper's posterior belief about Carrier  $i$ 's tendency to accept at  $(k_t^i, n_t^i)$  is

$$q_t^i \sim \text{Beta}(a^i + k_t^i, b^i + n_t^i).$$

Refer to  $(k_t^i, n_t^i)_{i=1,2,3}$  as a history and  $(a^i + k_t^i, b^i + n_t^i)_{i=1,2,3}$  as a belief state. Since the arms are independent, the shipper updates the belief state associated with Carrier  $i$  only when it is chosen.


Choosing Carrier  $i$  at belief state  $(a^i + k_t^i, b^i + n_t^i)$  has two sources of gains: the first is the expected instantaneous payoff, and the second is the informational value of learning about Carrier  $i$ . How the shipper optimally trades off these two sources of gains for a single-arm bandit problem at belief state  $(a, b)$  is summarized by the Gittins index, denoted by  $\text{Gittins}(a, b)$ . Under optimal learning, the shipper chooses

$$i^* = \arg \max_{i=1,2,3} \text{Gittins}(a^i + k_t^i, b^i + n_t^i) \quad \text{at history} \quad (k_t^i, n_t^i)_{i=1,2,3}.$$

**Gittins indices** The Gittins index in a single-arm bandit problem is the *stopping value*  $\lambda$  such that the shipper is indifferent between pulling the arm and stopping. Suppose that  $V_\lambda$  is the value function in the optimal stopping problem with stopping value  $\lambda$ . At belief state  $(a, b) \in (\mathbb{N}^+)^2$ ,

$$V_\lambda(a, b) = \lambda + \max \left\{ \frac{a}{b} - \lambda + \delta \left( \frac{a}{b} V_\lambda(a+1, b+1) + \left( 1 - \frac{a}{b} \right) V_\lambda(a, b+1) \right), 0 \right\}. \quad (15)$$

Figure 6: A simulated learning path



Notes: Priors:  $q_0^1 \sim \text{Beta}(12, 20)$ ,  $q_0^2 \sim \text{Beta}(8, 20)$ , and  $q_0^3 \sim \text{Beta}(5, 20)$ . True acceptance tendency:  $q^1 = 0.3$ ,  $q^2 = 0.25$ , and  $q^3 = 0.45$ . In this example, the shipper's prior is overly optimistic about Carriers 1 and 2, and overly pessimistic about Carrier 3. Thus, the Gittins indices of the former two carriers are generally decreasing, both because of the initial overoptimism and because the decrease in informational values as they are chosen. The shipper makes many switches between these two carriers before she starts to experiment with Carrier 3, at which time the Gittins index of Carrier 3 evolves.

Thus, for each  $\lambda$ , we can solve for  $V_\lambda$  as a fixed-point of the Bellman equation. The Gittins index at state  $(a, b)$  is the value of  $\lambda$  such that  $V_\lambda(a, b) = \lambda$ .

**A simulation exercise** Our simulation exercise involves two steps. In the first step, we estimate the Gittins index at each belief state, which involves the following sub-steps:

1. Let  $S = \{1, \dots, N\}^2$  for  $N = 200$  be the state space of beliefs and discretize  $[0, 1]$  for the value of the Gittins indices into a grid  $\Lambda$  with step size equal  $1/500$ .
2. For each value of  $\lambda \in \Lambda$ , find  $V_\lambda - \lambda$  via an iterative procedure using Equation (15). Notice that  $V_\lambda - \lambda$  is decreasing in  $\lambda$ .
3. For each  $(a, b) \in S$ , let  $\text{Gittins}(a, b) = \arg \min_{\lambda \in \Lambda} |V_\lambda(a, b) - \lambda|$ .

In the second step, we simulate a path of the shipper's choices and beliefs from  $t = 1$  to  $T = 500$ . In each period  $t = 1, 2, \dots$ , and with history  $(k_{t-1}^i, n_{t-1}^i)_{i=1,2,3}$ :

1. The shipper chooses  $i^* = \arg \max_{i=1,2,3} \text{Gittins}(a^i + k_{t-1}^i, b^i + n_{t-1}^i)$ .

2. Carrier  $i^*$  accepts with probability  $q^{i^*}$ , and the history evolves accordingly.

Figure 6 plots the evolution of the Gittins indices for a simulated history. Notice that the Gittins index of a carrier evolves only when that carrier is chosen. If learning steps are small, this feature results in the prevalence of the switch-back patterns. To see why, consider the first switch (at  $t$ ) from Carrier 1 to Carrier 2 and the next switch (at  $t'$ ) away from Carrier 2. With small learning steps, Carrier 1 would have a higher Gittins index than Carrier 3 at  $t$ . And since the Gittins indices of both Carriers 1 and 3 do not evolve between  $t$  and  $t'$ , their relative ranking remains the same until  $t'$ . This means that when switching away from Carrier 2, the shipper tends to switch back to Carrier 1.

## C Additional empirical evidence

**Relationship characteristics across carrier types** Our empirical evidence, as summarized by Table 5, shows substantial heterogeneity by carrier type in the nature of within-relationship interactions between shippers and carriers. Table 6 shows that key relationship characteristics determined at the formation stage of relationships also differ systematically by carrier type. First, large asset-based carriers have lower contract rates than small asset-based carriers and brokers, suggesting potential cost advantages of large asset-based carriers. Second, compared to an asset-based carrier, a broker tends to serve as the primary carrier for the same shipper on more lanes, but these lanes tend to have lower volume. This tendency helps explain why shippers use multi-lane punishments with brokers: First, low-volume lanes creates the need to pool incentives across lanes. Second, a shipper and broker interacting on many lanes means that the scope for incentive pooling is larger.

Table 6: Relationship characteristics by carrier type

	Large ABCs	Small ABCs	Brokers
Monthly lane volume	9.14	8.73	6.35
Contract rate - spot rate	-0.125	-0.0878	-0.0432
Inconsistency (loads / week)	1.04	1.12	1.17
Inconsistency (day of week)	0.615	0.546	0.555
Number of lanes	6.91	4.13	11.61

*Notes:* This table shows differences in relationship characteristics across carrier types. The first four rows give the means of four key shipper-carrier-lane characteristics. The last row indicates the average number of lanes comprising a shipper-carrier relationship.

**Role of relationship-specific investments** In the truckload freight setting, relationship-specific investments could take the form of a carrier adjusting his network of truck movements to best service a contracted lane. For example, a carrier can reduce empty miles by lining up a backhaul, thus making accepting forehaul loads less costly.<sup>68</sup> Among the three groups of carriers, we expect that such adjustments matter most for small ABCs; large ABCs have more flexibility in network adjustments due to their larger fleets, and brokers, who find carriers on the spot market, do not have similar concerns for empty miles because they do not incur the cost of such empty miles. Our analysis on carriers’ acceptance (Table 3) shows evidence consistent with this hypothesis. Only small ABCs have lower acceptance rates in the first seven days after being promoted from backup to primary status.

A breakdown of our endgame effect analysis (Section 5.1) by carrier type further supports the hypothesis that relationship-specific investments do not play an important role in the behavior of large ABCs and brokers. Figure 7 shows that following the announcement of RFP outcomes, the acceptance rate of both large ABCs and brokers, whether they won or lost the RFP, decreases significantly. This behavior cannot be driven by relationship-specific investments, including network adjustments. The reason is that winning carriers—who continue to be primary carrier the same lane—should have no incentives to undo their investments toward the end of the current contract period.

Table 7: Response of auction outcomes to carrier behavior

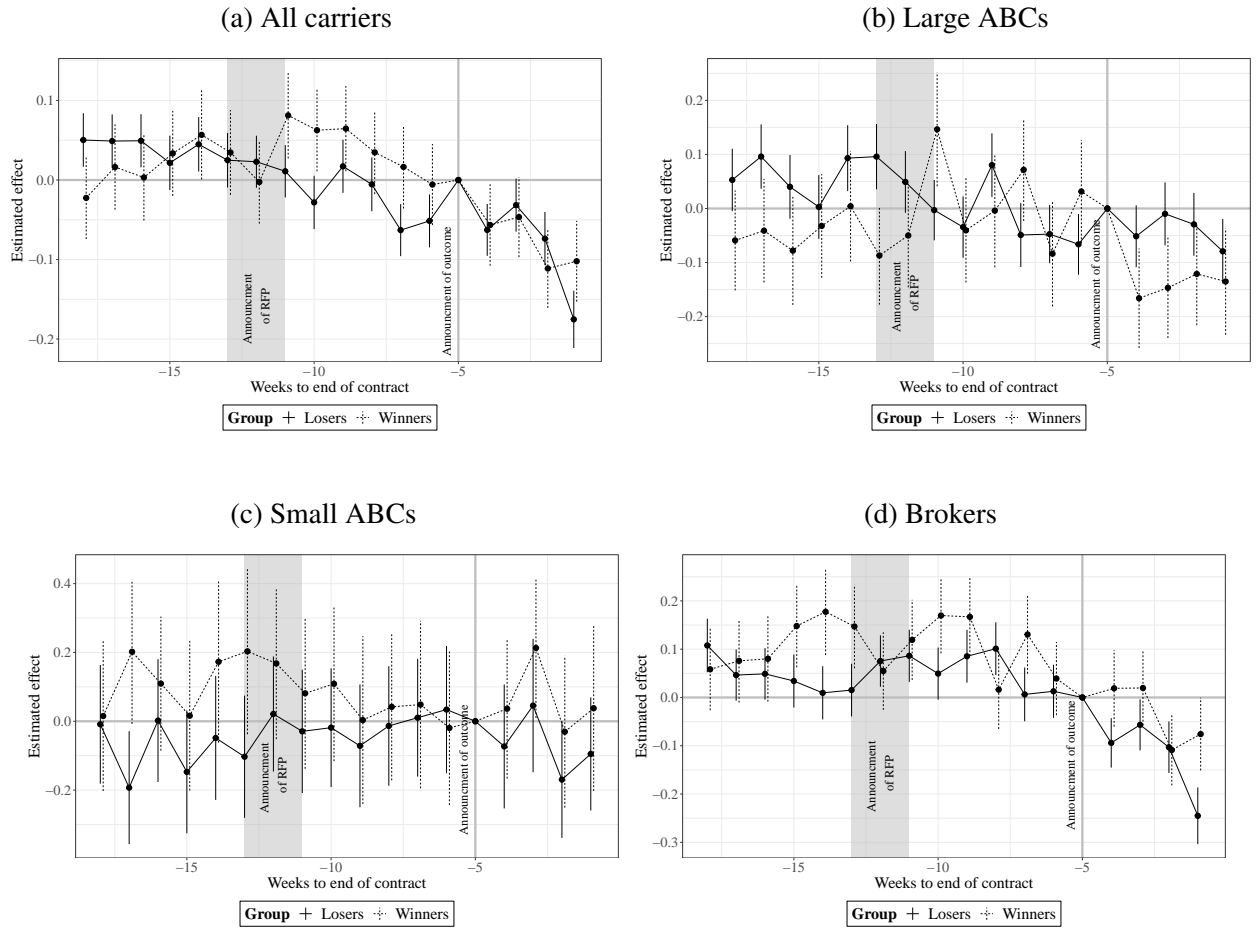
	Wins RFP	New contract premium
Rejection rate	-0.0676 (0.0239)	-0.430 (0.0798)
Observations	1673	373

*Notes:* Standard errors in parentheses. This table reports estimated coefficients for regressions of the form  $y_{scr}^{\ell} = \delta_0 + \delta_1 \text{Rejection rate}_{scr}^{\ell} + \delta_2 \text{Miles}^{\ell} + \epsilon_{scr}^{\ell}$  where  $\text{Rejection rate}_{scr}^{\ell}$  is the proportion of load offers rejected by carrier  $c$  from shipper  $s$  on lane  $\ell$  during the contract period and  $\text{Miles}^{\ell}$  is the length of lane  $\ell$ . In the first column, the outcome variable  $y_{scr}^{\ell}$  is an indicator for whether the primary carrier  $c$  wins RFP  $r$  and therefore maintains the primary position on lane  $\ell$ . In the second column, the outcome variable  $y_{scr}^{\ell}$  is the new contract rate of carrier  $c$  after the RFP (conditional on winning the RFP and maintaining the primary position). In keeping with the event study regressions above, both regression samples are limited to mass RFP events. The sample for the second column is further limited to primary carriers who win the RFP.

<sup>68</sup>The terminology of “forehaul” and “backhaul” is widely used in the industry. If a carrier has a contract for the lane from A to B (the forehaul), he would ideally also have a contract providing regular demand on the backhaul from B back to A.



Figure 7: End-of-contract effects by carrier type



Notes: Like Figure 4, this figure plots the estimated coefficients  $\{\alpha_k\}$  from equation (2) with the normalization  $\alpha_5 = 0$  is imposed. Panel (a) is same as Figure 4, while panels (b)-(d) present separate estimates by carrier type.

Table 8: Performance of promoted carrier relative to demoted carrier

	All types		Large ABCs	
	Acceptance	Contract rate	Acceptance	Contract rate
	-0.0122 (0.00244)	0.0911 (0.0275)	-0.150 (0.00654)	0.0448 (0.0813)
<i>N</i>	1009352	1009352	177030	177030

	Small ABCs		Brokers	
	Acceptance	Contract rate	Acceptance	Contract rate
	-0.184 (0.0105)	-0.185 (0.0903)	0.115 (0.00451)	-0.000292 (0.0159)
<i>N</i>	80686	80686	277878	277878

Notes: This table presents estimates from the following regression:

$$y_{sct}^{\ell} = \sum_{g \in \mathcal{G}} \beta_g \mathbb{I}\{g_{sct}^{\ell} = g\} + \underbrace{\gamma \left( \text{Spot rate}_t^{\ell} - \text{Contract rate}_{sct}^{\ell} \right)}_{\text{only included for } y=\text{Acceptance}} + \epsilon_{sct}^{\ell}$$

where  $\mathcal{G}$  represents a set of three groups of primary carriers: (a) RFP winners who are never demoted, (b) RFP winners who are eventually demoted, and (c) non-RFP winners who are promoted. The estimates listed in the table state the difference  $\beta_{(c)} - \beta_{(b)}$ . They therefore have the interpretation of how much the expectation of the outcome  $y$  changes when the shipper demotes the RFP winner and promotes his replacement. Outcomes are (1) an indicator for carrier  $c$  accepting load  $t$  and (2) carrier  $c$ 's contract rate. The results for all carrier types indicate that demoting the primary carrier tends to result a small (1.2pp) decrease in the acceptance rate while resulting in a sizeable (9 cent per mile) increase in the contract rate.