Long-Term Relationships and the Spot Market: Evidence from US Trucking*

Adam Harris[†] Thi Mai Anh Nguyen[‡]

August 26, 2024

Abstract

Long-term relationships play an important role in the economy, capitalizing on match-specific efficiency gains and mitigating incentive problems. However, the prevalence of long-term relationships can also lead to thinner, less efficient spot markets. We develop an empirical framework to quantify the market-level tradeoff between relationship and spot performance, and we apply this framework to the US truckload freight industry, one in which relationships with fixed-rate contracts predominate. We find that while the intrinsic benefits of relationships outweigh their negative externalities, social optimality requires a balance between relationship and spot transactions. The current institution comes reasonably close to achieving this balance, as the friction generated by the incompleteness of the current fixed-rate contracts acts as a partial corrective tax on relationship transactions. Overall, the current institution achieves 92% of the market-level first-best surplus, despite achieving only 64% of the relationship-level first-best surplus.

Key words: Long-Term Relationships, Spot Markets, Market-Thickness Externalities, Contracts, Trucking.

*We thank Nikhil Agarwal, Glenn Ellison, Bob Gibbons, Stephen Morris, Nancy Rose, Tobias Salz, and Mike Whinston for their guidance and support. We thank Angi Acocella, Chris Caplice, and many people in the truckload freight industry who generously shared their time, insights and data with us. We acknowledge financial support from the George and Obie Shultz Fund, the NSF GRFP (Harris), the Jerry Hausman Graduate Dissertation Fellowship (Nguyen), and the US DOT NSF grants SES1559013 and SES2315269 to NBER (Harris). The opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the US DOT and NBER, which provided funding, or of CH Robinson and DAT, which provided the data used in this research.

[†]Cornell University. Email address: adamharris380@gmail.com.

[‡]Corresponding author. New York University. Email address: anh.nguyentm01@gmail.com.

1 Introduction

Long-term relationships and spot transactions are ubiquitous—and typically competing—features of the economy (Baker et al., 2002). On the one hand, spot transactions present an outside option to relationships, making relational incentives harder to enforce; on the other hand, the formation and high performance of relationships could result in a thinner and less efficient spot market (Kranton, 1996). Given these counteracting forces, an equilibrium in which relationships predominate could result from either large intrinsic benefits of relationships or from a failure to coordinate to establish a thicker spot market.

We develop an empirical framework to quantify the market-level tradeoffs between long-term relationships and the spot market in US for-hire truckload freight. This is an economically important industry where 80% of all transactions occur within long-term fixed-rate contracts. Because these contracts are incomplete, they expose relationships to spot temptation; as a result, relationships achieve only 64% of the *relationship-level* first-best surplus. However, the benefits of any improvement in the performance of relationships would be offset by exacerbating the thinness and inefficiency of the spot market. We show that the current equilibrium strikes a reasonable balance between relationships' benefits and their externalities on the spot market, achieving 92% of the *market-level* first-best surplus.

The US for-hire truckload freight industry offers an ideal empirical setting for studying the interactions between relationships and the spot market. In this setting, a firm with demand for transportation service ("shipper") on an origin-destination pair ("lane") hires a trucking company ("carrier") for that service. Transactions occur either under long-term contracts that fix prices ("rates") but not volume or as spot arrangements.[1] The incompleteness of these long-term contracts allows for (i) rich dynamics within shipper-carrier relationships and (ii) an intricate interplay between these relationships and the spot market.

Two empirical advantages of our settings allow us to quantify both of these features. First, shippers use transportation management systems (TMS) to automate many aspects of their relationships with carriers. By utilizing the digital records of one TMS used by 51 shippers, we observe the details of all of these shippers' relationships, including their offers, carriers' responses, and each relationship's status at each point in time. Second, variation in demand results in substantial temporal and cross-lane variation in spot rates and volumes. Our framework exploits (i) temporal variation in spot rates to quantify relationship value and

[1] In the context of the international coffee market, Blouin and Macchiavello (2019) find that fixed-rate contracts expose relationships to strategic defaults but offer price insurance value. In the market for pulp, Tolvanen et al. (2021) show that buyers write quantity contracts to insure against trading frictions in the spot market. Our paper does not seek to explain why long-term contracts in our setting fix prices but not volumes. Instead, we seek to quantify the socially optimal split between relationship and spot transactions.

(ii) cross-lane variation in spot volumes to quantify the link between spot market thickness and efficiency.

In Section 3, we begin our analysis by establishing three data patterns that highlight the setting's key economic forces. First, one-third of the time, shippers do not select the carrier with the lowest contract rate as their "primary carrier". This suggests the importance of non-price factors in long-term relationships. Second, while carriers are more likely to turn down relationship transactions when spot rates are favorable to them, shippers are not. This suggests that the spot market crowds out relationships by creating a moral hazard problem only for carriers. We show in our companion paper Harris and Nguyen (2024) that shippers mitigate this moral hazard problem by conditioning relationship continuation on carrier acceptance. Third, consistent with Hubbard (2001), we find that spot arrangements take a larger share of total market volume on lanes with higher total demand. Scale efficiencies in the spot market offer one potential explanation for this pattern.

In Section 4, we develop a model that captures these patterns. At the relationship level, each relationship has two stages. In the first stage, the shipper's and carriers' match-specific gains are drawn from a distribution, and the shipper holds an auction to select one carrier with whom to form a relationship. In the second stage, the shipper and the winning carrier interact in a repeated game. In each period of this game, the shipper decides whether to terminate the relationship or to maintain the relationship and offer a load; the carrier decides whether to *accept* the offered load, reject it in favor of a *spot* offer, or reject it to remain *idle*. Delivering a load either within the relationship or on the spot market has some operational cost. On top of this, a relationship transaction realizes match-specific gains while a spot transaction incurs a search cost; the latter might scale (inversely) with spot market volume. At the market level, the spot market absorbs both direct spot demand and unfulfilled relationship demand, pinning down equilibrium spot rates and search costs.

In Section 5, we show that shippers' and carriers' primitives are identified from their actions in the auction and repeated game. To identify carriers' primitives, we exploit their optimal response to spot conditions in the repeated game. When spot rates are sufficiently low, a carrier accepts as long as his costs are low; when spot rates are sufficiently high, he never accepts. At the critical spot rate at which acceptance drops to zero, the carrier is indifferent between *accept* and *spot*. This indifference condition is key to separately identifying carriers' gains and costs. First, the critical spot rate pins down the carrier's relationship premium relative to spot transactions. Second, the carrier's decisions to the left of this critical point reveal information about his search and operational costs. Finally, to identify the causal link between search costs and spot market thickness, we use the predicted trade flows between US states from Caliendo et al. (2018) as a demand shifter.

To identify shippers' primitives, we exploit both carriers' bidding and the shipper's selection of the winning carrier in the auction stage. In equilibrium, the shipper selects the relationship with the highest total match-specific gains, and carriers effectively bid on how to split these gains into carriers' rents and shippers' rents. Under empirically plausible conditions, we show that there exists a monotone mapping between carriers' rents—whose components are either observed or already identified—and shippers' rents. We adapt techniques from Guerre et al. (2000) to pin down this mapping and recover shippers' rents. Conditioning these rents on contract rates yields shippers' match-specific gains.

In Section 6, we present our estimates of model primitives, showing that relationship transactions generate large gains to the participating parties but also exert substantial negative externalities on the spot market. Specifically, we find that the median relationship transaction generates \$0.78/mile for the shipper and \$0.68/mile for the carrier in relationship premiums. Savings on search costs account for more than a third of the total premium, and search costs account for about a third of carriers' total cost of servicing spot transactions. Crucially, we find that thicker spot markets substantially reduce search costs; doubling the spot market volume, for instance, would reduce search costs by \$0.36/mile.

In Section 7, we evaluate the current institution against alternatives that alter the role or performance of relationships. At the relationship level, we find that relationships generate substantial surplus to their participants. However, the current fixed-rate contracts capture only 64% of the relationship-level first-best surplus; adopting *index-priced* contracts that are pegged one-to-one to spot rates would achieve this first-best. Thus, our findings highlight the importance of relationships in the truckload setting while suggesting strong incentives for relationships' participants to switch to more flexible contracts.

At the market level, improving the performance of relationships comes at the cost of a thinner spot market with higher search costs. Higher search costs directly burden spot carriers, and stronger incentives to avoid search costs indirectly rationalize more relationship transactions with negative intrinsic gains or high operational costs. To evaluate this tradeoff, we derive the constrained first-best welfare for each level of the spot market's share of total market volume. Implementation involves combining relationship-optimal index-priced contracts with a corrective tax on relationship transactions. A zero tax would yield our *index-priced* counterfactual, while an infinite tax would yield our *spot-only* counterfactual.

We find that neither of these extreme counterfactuals achieves the market-level first-best welfare. Indeed, *spot-only* would result in substantial welfare loss relative to the current institution, suggesting that the dominance of relationships in the US truckload setting is not a failure to coordinate on forming a thick spot market. Furthermore, we find that the current institution realizes 92% of the market-level first-best surplus, with spot shares approximating

the first-best levels during periods of high demand. Intuitively, the incompleteness of the current fixed-rate contracts acts like a tax on relationship transactions that partially corrects the externalities of relationships on the spot market. This corrective mechanism works particularly well in periods of high demand, when carriers' spot temptations are high.

Related literature. Our paper contributes to three main strands of literature. First, we contribute to the literature on long-term relationships.[2] Conceptually, our paper is most closely related to Kranton (1996), who theorizes two-way crowding-out effects between long-term relationships and the spot market. Our contribution is an empirical framework that quantifies (i) both of these crowding-out effects and (ii) the socially optimal balance between relationship and spot transactions.

We build upon previous insights on how to quantify relationship primitives. Typical approaches rely on the conditions for dynamic enforcement (e.g., Startz, 2021) or optimal contracting (e.g., Perrigne and Vuong, 2011; Brugues, 2020) and exploit variation in spot market conditions to shift deviation profits (e.g., Macchiavello and Morjaria, 2015; Blouin and Macchiavello, 2019). We similarly exploit carriers' dynamic enforcement within relationships to recover their primitives. In doing so, however, we face two empirical challenges: on-path terminations, which leads to short panels (Kasahara and Shimotsu, 2009), and partially observed actions. We overcome these challenges by developing a dynamic discrete choice framework (Rust, 1994) and a novel support-based argument on spot rates. Another novelty of our approach is that we recover shipper primitives from the relationship formation process by formulating it as an auction on rents. This auction-based approach (Guerre et al., 2000) allows us to remain agnostic about the shipper's power to commit to a punishment scheme within her relationship with the contract-winning carrier.

The fact that the stickiness of long-term bilateral interactions could negatively affect overall market efficiency is not unique to our setting.[3] Zahur (2022) shows that formal long-term contracts for liquefied natural gas, while mitigating hold-up problems, hurt overall allocative efficiency by reducing firms' ability to respond to demand shocks. Our paper highlights an additional channel through which long-term relationships affect spot market performance: relationships increase spot market frictions by reducing its thickness, thereby

[2]This literature has established a rich set of empirical evidence on how long-term relationships generate value for participating parties and respond to external factors. For example, value creation in long-term relationships can arise from supply reliability (Adhvaryu et al., 2020; Cajal-Grossi et al., 2023), reputation building (Macchiavello and Morjaria, 2015), or relational adaptation (Barron et al., 2020). Relationships may terminate (Macchiavello and Morjaria, 2015) or restructure (Gil et al., 2021) in the face of large shocks, and can be hampered by competition (Macchiavello and Morjaria, 2021).

[3]More broadly, our paper relates to other empirical studies on the spillovers across coexisting modes of transaction. Such coexistence could generate quality dispersion (Galenianos and Gavazza, 2017), and market frictions could shift market modality (Gavazza, 2010, 2011; Startz, 2021).

reinforcing the appeal of relationships.[4]

Second, we contribute to the empirical literature on the efficiency of transportation markets. Previous work has shown that decentralized transactions generate search frictions, resulting in spatial and temporal misallocation (e.g., Lagos, 2003; Frechette et al., 2019; Buchholz, 2022). Large transportation platforms might mitigate this problem by exploiting economies of density and pricing instruments, but may also extract most of the surplus (e.g., Rosaia, 2023; Brancaccio et al., 2023); platform efficiency further depends on the sophistication of the pricing rule (e.g., Castillo, 2023) and the mechanism for eliciting users' heterogeneous preferences (e.g., Gaineddenova, 2022; Buchholz et al., 2024). In contrast to these papers, which study one-off transactions, our paper highlights the role of long-term bilateral relationships in an industry where such relationships predominate.

Third, we contribute to the literature on trucking. Early comparisons of long-term contracts and spot transactions in this industry showed that the former provide the benefit of savings on transaction costs (Masten, 2009) and that the latter are more popular on thicker lanes (Hubbard, 2001).[5] Our analysis connects these insights by showing that transaction costs in the spot market increase endogenously as more transactions occur in relationships rather than in the spot market. In quantifying the market-level tradeoffs between long-term relationship and spot transactions, we leverage our previous understanding of the scope and dynamic mechanisms within these relationships (Harris and Nguyen, 2024) and abstract from spatial allocation. Yang (2023) studies the home bias of truck drivers using a spatial equilibrium model but focuses exclusively on spot transactions.

2 Institutional details

The US for-hire truckload freight industry offers an ideal setting to study the market-level tradeoff between long-term relationships and the spot market. This section provides industry background and highlights two institutional features underlying this tradeoff: (i)

[4]This channel mirrors liquidity externalities in financial settings, where a natural solution is to centralize all transactions onto a single platform (e.g., Admati and Pfleiderer, 1988). In the market for Canadian government bonds, Allen and Wittwer (2023) show that market centralization is challenging because investors value relationships and dealer competition on the current platform is low.

[5]A related strand of literature studies asset ownership (Baker and Hubbard, 2003, 2004; Nickerson and Silverman, 2003) and the effects of deregulation on the trucking industry (Marcus, 1987; Rose, 1985, 1987; Ying, 1990). Since these earlier papers, significant improvements in how the trucking industry manages and tracks shipper-carrier interactions have generated richer transaction-level data that has been exploited in more recent studies. For example, the transportation and logistics literature has used such data to study the effects of long-term relationships on participating parties, examining reciprocity (Acocella et al., 2020), factors that affect carriers' value of relationships (Acocella et al., 2022b), and bilateral improvements from index pricing (Acocella et al., 2022a).

the incompleteness of long-term contracts and (ii) the fragmentation of spot transactions.

2.1 The US for-hire truckload freight industry

Trucking is the most important mode of transportation for US domestic freight. In 2019, trucks carried 72% of domestic shipments by value.[6] The US trucking industry comprises four segments: for-hire truckload, private truckload, less-than-truckload, and parcel.

In this paper, we focus on the largest of these segments: for-hire truckload. In this segment, a *shipper* (e.g., manufacturer, wholesaler, or retailer) with a *load* (shipment) to be transported on a *lane* (an origin-destination pair) on a specified date hires a *carrier* (i.e., trucking company) for that service.[7] As the name suggests, a *truckload* shipment fills an entire standard-sized truck. For-hire truckload carriers are therefore concerned about reducing miles traveled empty, and thus unpaid, but not about how to optimally combine shipments to fill up their trucks. The latter is a key concern of less-than-truckload and parcel carriers. This means that truckload carriers face simpler routing decisions and rely less on economies of scale, a difference that accounts for the framgentation of the truckload segment, in contrast to other segments (Ostria, 2003).[8] To further reduce heterogeneity, our analysis focuses on dry-van (as opposed to refrigerated, flatbed, or tanker) services and long-haul lanes of at least 250 miles.

Shippers and carriers in the US for-hire truckload industry engage in two main forms of transactions: transactions within long-term relationships and spot transactions. Long-term relationships predominate, accounting for 80% of total transacted volume. Spot arrangements account for the remaining 20% (Holm, 2020). The next two subsections provide institutional details for each of these market components.

2.2 Long-term relationships

Long-term relationships between shippers and carriers are formed via procurement auctions. Such an auction begins with a shipper asking for proposals from various carriers. Each carrier then submits a bid on a fixed contract rate to be charged on each load that the carrier transports for the shipper within the contract period (typically one or two years). Contract rates are accompanied by a fuel program, typically proposed by the shipper, that

[6]These statistics are calculated using data from the Bureau of Transportation Statistics.

[7]This is in contrast to private fleets, which are vertically integrated carriers serving a single shipper. Such vertical contractual arrangements tend to be chosen by companies that prioritize quality and reliability of service and those that have a dense network of truck movements that allows for efficient routing.

[8]Bokher (2018) reports that the top 50 truckload fleets account for only about 10% of the segment's total revenue, and about 90% of truckload fleets have fewer than six trucks.

compensates carriers for changes in fuel costs.[9] Based on these bids, the shipper then chooses a *primary carrier* and a set of backup carriers (who may receive requests for any loads that the primary carrier rejects).[10]

Having established the primary carrier and contract rates, the shipper uses a Transportation Management System (TMS) to automate her shipment requests to the carriers. When the shipper needs to transport a load on a lane, she inputs the details of the load into her TMS. Requests are then sent out to the carriers sequentially in the order of their ranks until a carrier accepts. Primary carriers are typically top-ranked, receiving all of the requests first, and are expected to accept most of the requests they receive.[11] Our analysis focuses on the relationships between shippers and their primary carriers.

A crucial and unique feature of this setting is that contracts between shippers and carriers are incomplete, fixing rates but not volume. On the one hand, carriers can reject a load requested by their contracted shippers without any legal recourse. On the other hand, shippers can influence the number of requests each carrier receives by changing their ranks in the TMS. That is, the initial ranking of carriers established in the auction is not permanent; rather, the shipper can *demote* a primary carrier to a lower rank, replacing him with a backup carrier at any time. Thus, the fact that contracts between shippers and carriers are incomplete leaves room both for potential opportunistic behaviors and for relational incentives to mitigate such opportunism.

2.3 Spot transactions

Spot transactions occur through multiple channels, all of which typically involve search and haggling. For example, a shipper or carrier can use an electronic load board either to post an available load or truck or to search existing posts. These load boards are marketplaces from which both sides can obtain contact information of potential matches, but rate negotiations are conducted offline. Shippers and carriers can also be matched either by brokers or by digital matching platforms, which employ real-time matching and pricing. For

[9] The most common fuel program calculates per-mile fuel surcharge as the per-mile difference between a fuel index and a peg, (index − peg)/escalator, where "escalator" (miles/gallon) is a measure of fuel efficiency. In practice, variation in the choice of the index, the peg and the escalator has little impact on shippers and carriers. For more details, see https://www.supplychainbrain.com/ext/resources.

[10] Typically, shippers ask for proposals on multiple lanes simultaneously and carriers are free to bid on a subset of them. See Caplice (2007) for more details on this procurement process.

[11] The process of sequential requests is sometimes referred to as a "waterfall" process. Most carriers take less than one hour to respond to a request, and the full waterfall process typically takes less than three hours. While backup carriers do not necessarily know their exact ranks, primary carriers need to know their (top) ranks so as to properly plan for future requests. Sometimes, due to specified capacity constraints of the primary carriers, requests are sent first to backup carriers. Table 4 in Appendix B provides an example of the sequential request process.

our purpose, we treat all of these channels as a single spot market with search or haggling costs that potentially vary with spot market thickness.

3 Data

3.1 Data description

We obtain detailed data on the interactions between shippers and carriers within long-term relationships. This data was generated by the transportation management system (TMS) software offered by TMC, a division of C.H. Robinson, which is the US's largest logistics firm. For each shipper in our data set and each lane of that shipper, we observe the details of all loads and requests. For each load, we observe the origin, destination, distance, and activity date.[12] For each request, we observe a timestamp, as well as the carrier's identity, contract rate, and accept/reject decision. We use the ordering of sequential requests for a given load to identify primary carriers and demotion events. Observed changes in carriers' contract rates allow us to identify auction events.[13]

Our TMS data spans the period from September 2015 to August 2019. In total, we observe 1.2 million loads and 2.3 million requests between 51 shippers and 1,933 carriers on more than sixteen thousand origin-destination pairs. Overall, 70.2% of loads are accepted by primary carriers, 19.7% by backup carriers, and 10.1% of shipments are rejected by all contracted carriers and thus fulfilled in the spot market. While carriers can either own assets or serve as brokers, our main analysis is restricted to the relationships between shippers and asset-owner carriers. Furthermore, we focus on primary carriers, their decisions and relationship status, treating loads fulfilled by backup carriers as spot arrangements.[14]

We obtain spot data from DAT Solutions, the dominant freight marketplace platform in the US and the leading vendor of spot market data. DAT divides the US into 135 Key Market Areas (KMAs). For each KMA-KMA lane, we observe weekly summary statistics for spot rates and spot volume. We merge this information with our TMS data using the TMS timestamps and origin-destination information. We exploit two kinds of variation in this spot rate data: temporal variation in spot rates enables us to identify the value of

[12] We also observe some performance measures, including whether carriers delivered on time and (for a subset of the data) the amount of time they spent at the destination unloading the shipment.

[13] See Appendix C for details of our data construction.

[14] Appendix F.1 compares the performance of asset-owners versus brokers and primary versus backup carriers on three dimensions: load acceptance, on-time delivery, and wait time at destination. Brokers accept more often but perform worse than asset-owners in other performance measures. This is because brokers, as intermediaries connecting shippers to multiple carriers, have different cost structure and relationships with shippers than do asset owners. Backup asset-owner carriers perform significantly worse than primary asset-owner carriers in all performance measures.

relationships, while cross-sectional variation in spot volume allows us to pin down the link between spot market thickness and efficiency.

On average, a relationship has 2.5 load requests per week, lasts for five months, and has 25% probability of ending in a demotion. Over our sample period, spot rates and contract rates are on average about \$1.85/mile, and both have large and persistent differences across lanes, with a spatial standard deviation (SD) of \$0.50/mile. However, there is substantially more temporal variation in spot rates (SD = \$0.23/mile) than in contract rates (SD = \$0.12/mile), reflecting the fact that contracts fix rates over an extended period of time. As a result, there are periods in which spot rates deviate substantially from contract rates, potentially tempting shippers and carriers to defect from their relationships. In recognition of these deviations between spot and contract rates over the course of the "freight cycle," our empirical analysis examines two separate market phases: the "tight" (high-demand) phase includes the years 2017 and 2018; the "soft" (low-demand) phase includes the remaining years. Figure 9 provides the time series of average contract rates and spot rates. Table 5 provides key summary statistics of our merged relationship and spot data.

3.2 Key facts

We begin by using the data to establish three empirical facts, which speak to three forces at the heart of the tradeoff between the performance of long-term relationships and the performance of the spot market: (i) the magnitude and heterogeneity of relationships' match-specific gains, (ii) the incentive problem within relationships, and (iii) the potential scale efficiencies in the spot market.

Fact 1. Relationships generate match-specific gains.

At the auction stage, the patterns of carriers' bidding and shippers' selection of primary carriers reflect substantial heterogeneity in how different shippers and carriers value potential relationships. Figure 1 plots the distribution of within-auction variation in contract rates (bids) across auctions. For each auction, the left panel shows the coefficient of variation of observed bids. The right panel shows the difference between the bids of the primary carrier and the lowest-bid backup carrier. In half of the auctions, the coefficient of variation of observed bids is at least 0.1, reflecting substantial differences in how different carriers value a potential relationship with the same shipper on the same lane.[15] Moreover, in one-third of the auctions, the primary carrier is not the lowest bidder, and in these instances, the median

[15]Since we do not observe the bids of low-ranked carriers who never receive shippers' requests, the observed price variation is a lower bound of the actual price variation. To mitigate this negative bias, Figure 1 excludes auctions with only one observed bid.

Figure 1: Within-auction variation in contract rates

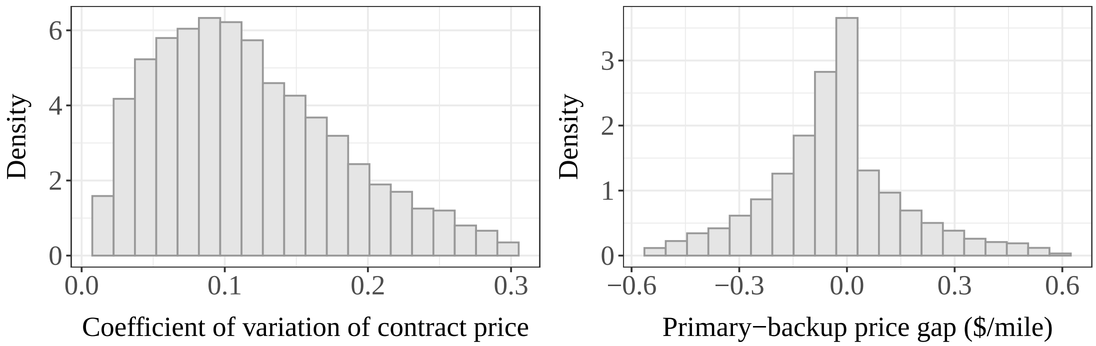

primary-backup price gap is \$0.16/mile. By revealed preference then, shippers care about non-price factors. Taken together, the two panels of Figure 1 show that, for both carriers and shippers, match-specific gains from relationships play a critical role.

These match-specific gains could arise from a number of different sources. Appendix F.1 provides evidence on two potential sources of match-specific gains, showing that primary carriers (i) are more likely to deliver goods on time and (ii) spend less time waiting at facilities relative to backup carriers.[16] The former may result from better alignment of a carrier's capacity (and other commitments) with a shipper's demand; the latter may result from better coordination between a shipper and a carrier when loading and docking. In addition, consistency in the timing of shipments helps carriers plan their networks of truck movements (Harris and Nguyen, 2024). Relationships can also differ in other dimensions, such as the insurance policies, billing practices, and en-route communication.

Fact 2. Carriers respond to spot temptation while shippers do not.

Since long-term contracts in our setting fix rates but not volumes, shippers and carriers might have an incentive to substitute spot transactions for relationship transactions when the former offer more favorable rates. The extent to which shippers and carriers respond to this incentive to deviate from relationships speaks to both the incentive problem created by fixed-rate contracts and the magnitude of shippers' and carriers' match-specific gains.

[16] On average, primary carriers are 6 pp more likely to be on-time (baseline of 66 pp) and spend 11 minutes less in wait time (baseline of 102 minutes) than backup carriers.

To assess carriers and shippers' response to spot temptation, we estimate two regressions:

$$\Pr(d_{ij\ell t} = \text{accept}) = \beta^c_{ij\ell,0} + \beta^c_1 \left| \frac{\tilde{p}_{\ell t} - p_{ij\ell}}{\tilde{\sigma}_\ell} \right|^+ + \epsilon^c_{ij\ell t}, \tag{1}$$

$$\ln(\text{Requests}^{\text{LT}}_{ij\ell,\text{month}}) = \beta^{\text{s}}_{ij\ell,0} + \beta^{\text{s}}_1 \left| \frac{\tilde{p}_{\ell,\text{month}} - p_{ij\ell}}{\tilde{\sigma}_\ell} \right|^- + \epsilon^s_{ij\ell,\text{month}}. \tag{2}$$

Here, i, j, and ℓ denote, respectively, a shipper, a primary carrier, and a lane. In (1), we estimate how carrier j's tendency to accept each request (t) responds to the premium of spot rate $\tilde{p}_{\ell t}$ over contract rate $p_{ij\ell}$, normalized by the standard deviation $\tilde{\sigma}_\ell$. Likewise, in (2), we estimate how the number of loads shipper i offers to the primary carrier j in a given month responds to the median spot premium in that month. Both regressions include relationship fixed effects to control for heterogeneity across relationships. Moreover, the carrier (shipper) regression focuses on instances when spot rates are higher (lower) than contract rates, as this is when the temptation to deviate to the spot market is highest. If shippers and carriers respond opportunistically to spot conditions, then $\beta^c_1 < 0$ and $\beta^s_1 > 0$.

Table 1 presents our estimates of regressions (1) and (2). As contract rates may be correlated with unobserved heterogeneity not captured by fixed effects, our (IV) specification instruments for spot premiums using the gaps between the current spot rates and the spot rates at the time of the auction. Our (IV) estimates show that carriers respond to spot rates while shippers do not. A one-standard deviation increase in the spot premium decreases carrier acceptance by 8.7 pp.

Table 1: Shippers and carriers' response to spot temptation

Response to spot rate	Expected sign	(OLS)	(IV)
Carrier's acceptance (β^c_1)	$-$	-0.070	-0.087
		(0.001)	(0.001)
Shipper's volume (β^s_1)	$+$	-0.020	-0.014
		(0.003)	(0.008)

These findings suggest that the incentive problem in long-term relationships is manifest in carriers rejecting shippers' requests when spot rates are high. As argued in our companion paper Harris and Nguyen (2024), carriers' incentive problem is mitigated directly by carriers' match-specific gains and indirectly by shippers' punishment scheme. Specifically, shippers condition carriers' primary status on carriers' past rejections, punishing their rejections today by denying them future requests. We find that the incentives this scheme creates are economically meaningful, but the scheme is soft. Thus, the limited response of carriers to spot temptation suggests that carriers enjoy meaningful gains from relationships. Shippers

may enjoy even larger such gains, given their lack of response to spot temptation.[17]

Fact 3. Spot transactions take a larger share on thicker lanes.

Finally, we present suggestive evidence of scale efficiencies in the spot market. Intuitively, if the spot market becomes more efficient as it becomes thicker, then there is an equilibrium force that—all else equal—results in a higher share of spot transactions on lanes with higher potential for total market volume. To test this hypothesis, we estimate the following cross-sectional regression:

$$\ln(\text{Volume}^{\text{spot}}_{ss'}/\text{Volume}^{LT}_{ss'}) = \beta_0 + \beta_1 \ln(\text{Volume}^{\text{total}}_{ss'}) + \text{controls}_{ss'} + \epsilon_{ss'}. \tag{3}$$

On each state-to-state lane, $\text{Volume}^{\text{spot}}_{ss'}$ is the average weekly load posts in the spot market, $\text{Volume}^{\text{LT}}_{ss'}$ is the average weekly loads accepted within long-term relationships, and $\text{Volume}^{\text{total}}_{ss'}$ is the total for-hire truckload volume in 2017 taken from the Commodity Flow Survey. We run a log-regression to avoid scaling issues between different data sets. An estimate of $\beta_1 > 0$ would be consistent with our hypothesis that the spot market's share of total market volume increases with total market volume.

Table 2 presents our estimates of Equation (3). Compared to the baseline specification (OLS1), (OLS2) includes controls for demand characteristics. The last two specifications further isolate unobserved cost factors by instrumenting for total volume with demand shifters. (IV1) uses the predicted trade flows across different states of the US (Caliendo et al., 2018), and (IV2) uses origin and destination population densities.[18] Our preferred specification (IV1) suggests that doubling market thickness could increase spot share on a lane from 20% to 33%. This large quantitative effect suggests a potentially strong link between spot market thickness and efficiency.

[17] In contrast, Casaburi et al. (2019) document opportunism on the buyer's side in Kenyan farming contracts.

[18] Demand characteristics included in (OLS2) are the frequency and timing consistency of shipments. As expected, more frequent and consistent demand is more desirable to carriers, favoring relationship over spot transactions. (IV1) proxies for demand on state-to-state lanes by the calibrated trade flows in Caliendo et al. (2018). Their model captures input-output linkages between difference sectors, labor mobility, and heterogeneous productivities; they calibrate this model to 2012 data aggregating trades across all modes of transportation. Since there might still be endogeneity concerns with using calibrated trade flows, we include for robustness (IV2), which proxies demand simply by origin and destination densities. This latter method suggests an even stronger connection between market thickness and the desirability of spot transactions.

Table 2: Suggestive evidence of scale efficiency in the spot market

$\ln(\text{Volume}^{\text{spot}}/\text{Volume}^{\text{LT}})$	(OLS1)	(OLS2)	(IV1)	(IV2)
$\ln(\text{Volume}^{\text{total}})$	−0.0608	0.0324	0.130	0.226
	(0.036)	(0.032)	(0.040)	(0.068)
ln(distance)	−0.249	−0.297	−0.197	−0.0994
	(0.075)	(0.065)	(0.070)	(0.090)
Frequency		−0.0821	−0.0803	−0.0786
		(0.013)	(0.013)	(0.013)
Inconsistency		1.176	1.213	1.249
		(0.071)	(0.071)	(0.075)
Instruments				
Predicted trade (Caliendo et al., 2018)			✓	
Origin & destination densities				✓

Note: Frequency is the median average monthly volume of a shipper on a lane; Inconsistency is the median coefficient of variation of a shipper's loads in a week over the four weeks of a month. These regressions aggregate observed spot and long-term relationship volumes to the state-to-state level, resulting in 1250 state-to-state observations. All specifications include indicators of a lane's origin or destination being in the Midwest to control for the Midwest's overrepresentation in our microdata.

4 Model

Quantifying the market-level value of long-term relationships requires a model that captures (i) the value created within relationships and (ii) the interactions between these relationships and the spot market. Figure 2 provides an overview of our model.

At the relationship level, a shipper selects a primary carrier via an auction and interacts with this carrier in a repeated game under the fixed contract rate established in the auction. In each period of this game, the shipper decides whether to terminate the relationship or maintain it and offer the carrier a load; if a load is offered, the carrier decides whether to accept or reject it. The premium of a relationship transaction over a spot transaction comes not only from the match-specific gains to the shipper and carrier, but also from savings on search costs for spot transactions. At the market level, long-term relationships and the spot market interact in two ways. First, spot temptation sometimes leads carriers to reject load requests within relationships. Such rejections create an overflow of demand and supply from relationships to the spot market, mediating spot rates. Second, spot volume affects search costs and thus the relative desirability of spot transactions.

Figure 2: Model overview

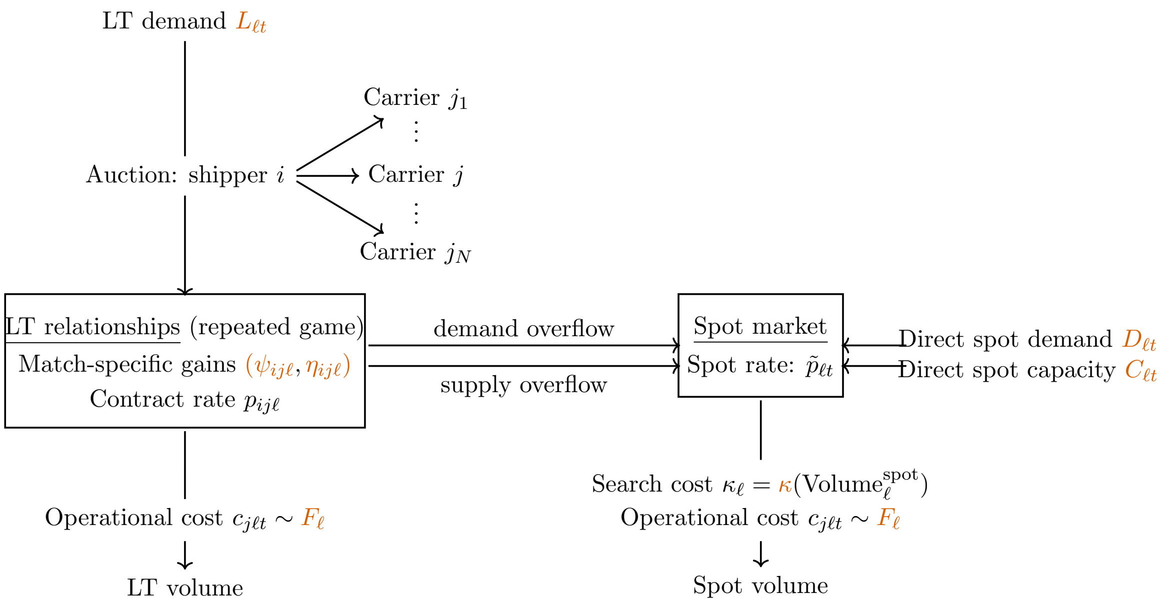

4.1 Individual relationship

4.1.1 Model primitives

Since the for-hire truckload market is very competitive, with many shippers and carriers, we assume that individual shippers and carriers perceive the market process as exogenous. In particular, these agents perceive the spot rate $\tilde{p}_{\ell t}$ on lane ℓ to follow an AR(1) process $\mathcal{P}_\ell$, and the search cost for spot transactions κ_ℓ on lane ℓ to vary with the average spot volume $\text{Volume}_\ell^{\text{spot}}$ according to some function κ.

Per-period payoffs. The relationship between a shipper i and a carrier j on lane ℓ is characterized by a tuple $(\psi_{ij\ell}, \eta_{ij\ell}, p_{ij\ell}, \delta_{i\ell})$ of relationship characteristics and a tuple $(\mathcal{P}_\ell, F_\ell, \kappa_\ell)$ of lane characteristics. Here, $\psi_{ij\ell}$ is the shipper's match-specific gain from transacting with the carrier on lane ℓ; $\eta_{ij\ell}$ is the carrier's match-specific gain from transacting with the shipper on lane ℓ; $p_{ij\ell}$ is the contract rate; $\delta_{i\ell}$ is the discount factor, reflecting the frequency of interactions; $\mathcal{P}_\ell$ is the spot process; F_ℓ is the distribution of the carrier's operational costs on lane ℓ; and $\kappa_\ell = \kappa(\text{Volume}_\ell^{\text{spot}})$ is the carrier's cost of searching for a spot load on lane ℓ, which depends on the average spot volume on that lane. Denote by $\tilde{p}_{\ell t}$ and $c_{j\ell t}$, respectively, the spot rate and the operational cost of the carrier on lane ℓ in period t. The carrier gets a period-t payoff of $v_{ij\ell t} = \eta_{ij\ell} + p_{ij\ell} - c_{j\ell t}$ when servicing the contracted shipper's load and $v_{ij\ell t} = \tilde{p}_{\ell t} - \kappa_\ell - c_{j\ell t}$ when servicing a spot load; otherwise, $v_{ij\ell t} = 0$. The shipper gets a period-t payoff of $u_{ij\ell t} = \psi_{ij\ell} - p_{ij\ell}$ if her load is serviced by the contracted carrier; else

the shipper gets service from the spot market and pays the spot rate, $u_{ij\ell t} = -\tilde{p}_{\ell t}$. Thus, relative to a spot transaction, a relationship transaction yields non-price premiums of $\psi_{ij\ell}$ to the shipper and $\eta_{ij\ell} + \kappa_\ell$ to the carrier.

Actions and timing. Each relationship between a shipper i and carrier j on lane ℓ comprises a formation stage ($t = 0$) and a repeated game ($t \geq 1$). In each period of the repeated game, the shipper decides between terminating the relationship and offering a load to the carrier; if the shipper offers a load, the carrier decides whether to accept or reject it. Denote by $d_t \in \{\text{accept, spot, idle}\}$ the carrier's decision in period t: $d_t = \text{accept}$ if the carrier accepts the offered load; $d_t = \text{spot}$ if the carrier rejects the offered load to serve the spot market; and $d_t = \text{idle}$ if the carrier rejects and remains idle.[19] We allow the shipper to condition relationship termination on an index that summarizes the carrier's past rejections, defined recursively by $I_R : (R_{t-1}, d_t) \mapsto R_t = \alpha R_{t-1} + (1-\alpha)\mathbf{1}\{d_t \neq \text{accept}\}$. The weight α and the initial state R_0 are known.

In the auction stage, shipper i holds an auction a to select a primary carrier (or none). The timing of this stage is as follows:

- Shipper i announces the expected frequency of interactions $\delta_{i\ell}$, other lane characteristics, and a punishment scheme $\sigma_0 : (R_{t-1}, \tilde{p}_t) \mapsto [0, 1]$, which specifies for each rejection index and current spot rate the probability that the shipper would demote a primary carrier in the repeated game.

- A set J_a of N carriers arrive. For each $j' \in J_a$, a pair of shipper-carrier match-specific gains $(\psi_{ij'\ell}, \eta_{ij'\ell})$ are drawn i.i.d. from a distribution $G_\ell^{\psi,\eta}$. The shipper's match-specific gain $\psi_{ij'\ell}$ is observable both to her and to carrier j', while the carrier's match-specific gain $\eta_{ij'\ell}$ is privately known to carrier j' alone.[20]

- Each carrier j' proposes a contract rate $p_{ij'\ell}$.

- Shipper i chooses either to form a relationship with a carrier $j \in J_a$ or to always go to the spot market.

If shipper i chooses j as the primary carrier, then the shipper and carrier interact repeatedly until the carrier is demoted. Starting in period $t \geq 1$ with public Markov state $(R_{t-1}, \tilde{p}_{\ell t})$, the relationship evolves to the next period as follows:

[19] Here, *idle* captures all of the carrier's outside options, including the opportunity to service on other lanes.

[20] This informational assumption matches certain features of the communication process between shippers and carriers. When asking for proposals, a shipper details her preferences for the service on a lane, and carriers respond with proposals explaining how they can meet such preferences. It is harder for the shipper to know how much carriers value their relationships, since this further depends on carriers' internal operations.

– Nature draws $\tilde{p}_{\ell t} \sim \mathcal{P}_\ell(\cdot|\tilde{p}_{\ell t-1})$ and $c_{j\ell t} \sim \mathcal{F}_\ell$. The spot rate is observed by both the shipper and carrier; the operational cost is privately observed by the carrier.

– The shipper decides whether to demote the primary carrier according to the probabilistic punishment scheme σ_0 announced in the auction stage.

– If the carrier is demoted, both the shipper and carrier resort to the spot market for all future transactions. If the relationship is maintained, the shipper offers a load to the carrier. The carrier responds to this offer according to a Markov strategy $\sigma_{ij\ell} : (R_{t-1}, \tilde{p}_{\ell t}) \to \{\text{accept, spot, idle}\}$.

– The Markov state evolves to $(\alpha R_{t-1}, \tilde{p}_{\ell t})$ if the carrier accepts the load offered within the relationship and $(\alpha R_{t-1} + (1-\alpha), \tilde{p}_{\ell t})$ otherwise.

Two modeling assumptions help us discipline the equilibrium behaviors of shippers and carriers in long-term relationships. First, carriers bidding on the same lane have the same cost distribution and search costs but differ in the match-specific gains $(\psi_{ij\ell}, \eta_{ij\ell})$ that they would generate in a relationship with the shipper. This means that, across potential relationships within an auction, the shipper's and carriers' expected payoffs depend only on the resulting per-transaction rents: $(\psi_{ij\ell} - p_{ij\ell})$ for the shipper and $(\eta_{ij\ell} + p_{ij\ell})$ for the winning carrier. Second, since the shipper's punishment scheme does not condition on carriers' identities or bidding behaviors, each carrier's expected payoff depends only on his own per-transaction rent. On the other hand, the shipper's expected payoff depends both on her own per-transaction rent and on the carrier's per-transaction rent. This is because the shipper's rent is realized only when the carrier accepts, and acceptance is more likely when the carrier's rent is higher.

For a relationship with per-transaction rents $(\psi_{ij\ell} - p_{ij\ell}, \eta_{ij\ell} + p_{ij\ell})$, denote by $V(R_{t-1}, \tilde{p}_{\ell t-1}|\eta_{ij\ell} + p_{ij\ell})$ the carrier's expected payoff and $U(R_{t-1}, \tilde{p}_{\ell t-1}|\psi_{ij\ell} - p_{ij\ell}, \eta_{ij\ell} + p_{ij\ell})$ the shipper's expected payoff in period t conditional on the rejection index and spot rate at the end of the last period. Write $\underline{V}(\tilde{p}_{\ell t})$ and $\underline{U}(\tilde{p}_{\ell t})$ for the termination payoffs of the carrier and shipper, given the current spot rate $\tilde{p}_{\ell t}$.

4.1.2 Equilibrium behaviors

We derive the equilibrium behaviors of individual shippers and carriers by working backwards. First, we derive the optimal dynamic play of primary carriers in the repeated game. This yields the expected payoffs of a shipper and a carrier in a relationship as functions of their per-transaction rents. Then, taking these expected payoffs as input, we derive carriers' optimal bidding and shippers' optimal selection of primary carriers in the auction.

Repeated game. Consider a primary carrier with rent $\eta_{ij\ell}+p_{ij\ell}$. The optimal play of this carrier at state $(R_{t-1}, \tilde{p}_{\ell t})$ depends on his *full compensation*, or how much he is compensated for an acceptance,

$$\bar{p}(R_{t-1}, \tilde{p}_{\ell t}|\eta_{ij\ell}+p_{ij\ell}) \equiv \underbrace{\eta_{ij\ell}+p_{ij\ell}+\kappa_\ell}_{\text{relationship premium}} + \frac{\delta}{1-\delta}\underbrace{\left(V(\alpha R_{t-1}, \tilde{p}_{\ell t}|\eta_{ij\ell}+p_{ij\ell}) - V(\alpha R_{t-1}+(1-\alpha), \tilde{p}_{\ell t}|\eta_{ij\ell}+p_{ij\ell})\right)}_{\text{dynamic compensation}}. \tag{4}$$

The first component in this expression—the relationship premium—is the carrier's direct benefit from a relationship transaction relative to a spot transaction, including his per-transaction rent and his savings on search costs. The second component—the dynamic compensation—is the difference in the carrier's continuation payoff were he to accept versus were he to reject the offered load. This latter component is positive if the shipper punishes a higher rejection rate with a higher likelihood of demotion.

The carrier's optimal decision reduces to choosing the maximum among the following:

$$\begin{cases} \bar{p}(R_{t-1}, \tilde{p}_{\ell t}|\eta_{ij\ell}+p_{ij\ell}) - \tilde{c}_{j\ell t} & \text{, if } d_t = \text{accept} \\ \tilde{p}_{\ell t} - \tilde{c}_{j\ell t} & \text{, if } d_t = \text{spot} \\ 0 & \text{, if } d_t = \text{idle,} \end{cases} \tag{5}$$

where $\tilde{c}_{j\ell t} \equiv c_{j\ell t}+\kappa_\ell$ is the *transformed cost*, or the carrier's cost of servicing a spot load. Let the optimal strategy $\sigma_{ij\ell} : (R_{t-1}, \tilde{p}_{\ell t}) \mapsto \{\text{accept, spot, idle}\}$ of this carrier be the solution of (5) that breaks ties in favor of acceptance. Together, the shipper's punishment scheme σ_0 and the carrier's optimal strategy $\sigma_{ij\ell}$ dictate the evolution and expected payoffs of each relationship.

Auction. We focus on the class of *symmetric monotone equilibria* for carriers' bidding and shippers' selection of primary carriers. The key advantage of this class of equilibria is that they reduce the auction in our setting to an auxiliary first-price auction in which carriers propose *effective bids* and the shipper selects the carrier with the highest such bid.

Such a transformation is non-trivial for two reasons. First, our model has two-sided match-specific gains. Second, the shipper (i.e., the auctioneer) prefers higher rents not only for herself but also for the winning carrier. We argue that both issues can be resolved by combining our informational assumptions on the match-specific gains $(\psi_{ij\ell}, \eta_{ij\ell})_{j\in J_a}$ with two equilibrium selection conditions. Letting $\theta_{ij\ell} \equiv \psi_{ij\ell}+\eta_{ij\ell} \in \Theta$ denote the *total match quality* between shipper i and carrier j on lane ℓ,

Condition 1 (Single indexing) The equilibrium (per-transaction) rents of the shipper and carrier are $\psi_{ij\ell} - p_{ij\ell} = \mathbf{b}(\theta_{ij\ell})$ and $\eta_{ij\ell} + p_{ij\ell} = \theta_{ij\ell} - \mathbf{b}(\theta_{ij\ell})$, respectively, for some function $\mathbf{b}$.

Condition 2 (Monotonicity) The shipper's rent $\mathbf{b}(\theta_{ij\ell})$ and carrier's rent $\theta_{ij\ell} - \mathbf{b}(\theta_{ij\ell})$ are both strictly increasing in $\theta_{ij\ell}$.

Condition 1 says that in equilibrium, match-specific gains should affect which relationship is formed and the resulting per-transaction rents only via the total match quality. That is, the split of this total match quality into the shipper's and the carrier's per-transaction rents can depend on the competitiveness of the auction and both sides' outside options, but not how much each side contributes to the total match quality. Condition 2 says that in equilibrium, both the shipper and the carrier strictly benefit from higher total match quality.[21]

We now formulate the auxiliary auction. First, recall our informational assumption that match-specific gains $(\psi_{ij\ell}, \eta_{ij\ell})$ are drawn i.i.d. across carriers, with $\psi_{ij\ell}$ being known by shipper i and carrier j and $\eta_{ij\ell}$ privately known by carrier j. This assumption implies that (i) total match qualities $\theta_{ij\ell}$ are i.i.d. across carriers, and (ii) when bidding a contract rate $p_{ij\ell}$, the carrier effectively chooses the shipper's per-transaction rent, $\mathbf{b}(\theta_{ij\ell}) = \psi_{ij\ell} - p_{ij\ell}$. Second, under Condition 2, the shipper can simply select the carrier that proposes the highest shipper rent, as that carrier also secures the highest carrier rent and will thus accept most frequently among all carriers. In summary, the auxiliary auction has an i.i.d., one-dimensional, type $\theta_{ij\ell}$; carriers bid on the implied shipper rents $\mathbf{b}(\theta_{ij\ell})$, and the shipper selects the carrier that proposes the highest shipper rent.

Assume that the outside option for both the shipper and carrier is having access only to the spot market. We define symmetric monotone equilibria formally as follows.

Definition 1. (Symmetric monotone equilibrium) In a symmetric monotone equilibrium, the following conditions hold:

(i) There exists a differentiable function $\mathbf{b} : \Theta \to \mathbb{R}$ that satisfies Conditions 1 and 2. Furthermore, carriers choose their effective bids optimally for all $\theta_{ij\ell}$,

$$\mathbf{b}(\theta_{ij\ell}) = \arg\max_{b} G^{\theta}_{\ell}(\mathbf{b}^{-1}(b))^{N-1}(V(R_0, \tilde{p}_{\ell 0}|\theta_{ij\ell} - b) - \mathbf{E}[\underline{V}(\tilde{p}_{\ell 1})|\tilde{p}_{\ell 0}]).$$

The shipper chooses the carrier j with the highest effective bid subject to the shipper's

[21] These conditions are analogous to those of Nash bargaining (Nash, 1953). First, both parties' payoffs depend only on their joint surplus, outside options, and bargaining power. That is, contract prices in our setting are analogous to transfers in Nash bargaining models; in relationships with the same total match quality but different match-specific gains, equilibrium contract prices move to maintain the same split of this match quality into per-transaction rents. Second, as long as both parties have positive bargaining power, their payoffs under Nash bargaining are strictly increasing in the joint surplus.

expected payoff from that relationship being no less than her outside option,

$$j \in \arg\max_{j' \in J_a} \mathbf{b}(\theta_{ij'\ell}) \text{ s.t. } U(R_0, \tilde{p}_{\ell 0} | \mathbf{b}(\theta_{ij\ell}), \theta_{ij\ell} - \mathbf{b}(\theta_{ij\ell})) \geq \mathbf{E}[\underline{U}(\tilde{p}_{\ell 1})|\tilde{p}_{\ell 0}].$$

(ii) If the auction selects j as the primary carrier, then in the repeated game, j uses the optimal strategy $\sigma_{ij\ell} : (R_{t-1}, \tilde{p}_{\ell t}) \mapsto \{\text{accept, spot, idle}\}$ associated with rent level $\theta_{ij\ell} - \mathbf{b}(\theta_{ij\ell})$.

4.2 Market equilibrium

Denote by $L_{\ell t}$ the measure of shippers who establish long-term relationships on lane ℓ in period t, by $D_{\ell t}$ the measure of direct spot demand by shippers who do not establish relationships, and by $C_{\ell t}$ the capacity of spot carriers. Let $\mu_{\ell t}(.|\tilde{p}_{\ell t})$ denote the distribution of full compensations $\bar{p}$ over all relationships. The status of each relationship, as captured by the rejection index, is embedded in the measure $\mu_{\ell t}(.|\tilde{p}_{\ell t})$. Higher $\mu_{\ell t}(.|\tilde{p}_{\ell t})$ in the FOSD-sense means higher aggregate acceptance probability.

The market equilibrium condition is

$$L_{\ell t} + D_{\ell t} = \underbrace{L_{\ell t} \int_{\tilde{p}_{\ell t}}^{\infty} F(\bar{p} - \kappa_\ell) d\mu_{\ell t}(\bar{p}|\tilde{p}_{\ell t})}_{\text{LT-relationship volume}} + \underbrace{[L_{\ell t}\mu_{\ell t}(\tilde{p}_{\ell t}|\tilde{p}_{\ell t}) + C_{\ell t}] F(\tilde{p}_{\ell t} - \kappa_\ell)}_{\text{spot volume}}. \tag{6}$$

Within long-term relationships, a shipper's requests are accepted if the full compensation $\bar{p}$ of the carrier is higher than the current spot rate and higher than the sum of the carrier's search and operational costs. Loads offered but rejected in long-term relationships will be fulfilled in the spot market, either by carriers in relationships with low full compensations or by those not in relationships.[22] A positive aggregate demand shock, either from an increase in demand for long-term relationships or from an increase in direct spot demand, increases the equilibrium spot rate.[23]

For the welfare analysis of different market institutions, we keep fixed the long-term demand $L_{\ell t}$, the direct spot demand $D_{\ell t}$, and the total capacity of carriers, including the spot capacity $C_{\ell t}$ and that of those who form relationships $L_{\ell t}$. Additionally, we normalize the value of spot interactions for shippers who want to establish long-term relationships to

[22]For simplicity, we count backup carriers as spot carriers.

[23]While we do not provide conditions on the fundamental process of shipping demand and carrier capacity that would give rise to an AR(1) process of spot rates, note that our model and identification argument could accommodate any autocorrelated process of spot rates. The restriction to AR(1) spot process is to facilitate estimation.

zero. Thus, the market-level welfare of the current institution on lane ℓ in period t is

$$\begin{aligned}W^0_{\ell t} =& L_{\ell t}\int_{\tilde{p}_{\ell t}}^{\infty}(\mathbf{E}[\theta|\bar{p},\tilde{p}_{\ell t}]-\mathbf{E}[c_{\ell t}|c_{\ell t}\leq\bar{p}-\kappa_\ell])F(\bar{p}-\kappa_\ell)d\mu_{\ell t}(\bar{p}|\tilde{p}_{\ell t})\\&+[L_{\ell t}\mu_{\ell t}(\tilde{p}_{\ell t}|\tilde{p}_{\ell t})+C_{\ell t}](-\kappa_\ell-\mathbf{E}[c_{\ell t}|c_{\ell t}\leq\tilde{p}_{\ell t}-\kappa_\ell])F(\tilde{p}_{\ell t}-\kappa_\ell).\end{aligned}$$

Alternative institutions that change the share or dynamics of long-term relationships modify the measure $\mu_{\ell t}(.|\tilde{p}_{\ell t})$ and the conditional expected match quality $\mathbf{E}[\theta|\bar{p},\tilde{p}_{\ell t}]$.

5 Identification

This section provides intuition for our identification argument, which is formalized in Theorem 1. We identify model primitives sequentially, focusing on the repeated game for carrier primitives and on the auction for shipper primitives.

Suppose that in each relationship, we observe the contract rate $p_{ij\ell}$, the duration $T_{ij\ell}$, and for each period $t \leq T_{ij\ell}$, the spot rate $\tilde{p}_{\ell t}$ and whether the carrier accepts ($d_{ij\ell t} = \text{accept}$) or rejects ($d_{ij\ell t} \in \{\text{spot}, \text{idle}\}$). Furthermore, suppose that we observe the number n_a of bidders in each auction who pass the shipper's individual rationality constraint.[24] To simplify notation, assume that the discount factor is $\delta_{i\ell} = \delta$. Objects that we do not observe and want to identify are the lane-specific distribution F_ℓ of operational costs, search cost κ_ℓ, the shipper's punishment scheme σ_0, and the distribution $G^{\psi,\eta}_\ell$ of match-specific gains $(\psi_{ij\ell}, \eta_{ij\ell})$ across relationships. Our identification argument relies on the following assumptions, which will be maintained throughout our analysis.

Assumption 1. *(Regularity) Assume that the following regularity conditions hold:*

(i) The AR(1) spot process $\mathcal{P}_\ell$ of $\tilde{p}_{\ell t}$ conditional on $\tilde{p}_{\ell t-1}$ has support $\mathbb{R}^+$, $\forall \tilde{p}_{\ell t-1}$.

(ii) The shipper's strategy satisfies $\sigma_0(R_{t-1}, \tilde{p}_{\ell t}) < 1, \forall (R_{t-1}, \tilde{p}_{\ell t})$.

(iii) The underlying distribution of match-specific gains $G^{\psi,\eta}_\ell$ has full support in $\mathbb{R}^2$. Moreover, it induces an underlying distribution of match quality $G^\theta_\ell \equiv G^{\psi+\eta}_\ell$ that has a strictly decreasing hazard rate, $g^\theta_\ell / G^\theta_\ell$.[25]

(iv) The distribution F_ℓ of operational costs is Normal(μ^c_ℓ, σ^c).

(v) There is a demand shifter z_ℓ independent of operational costs.

Assumption 2. *(Properties of full compensation schedules) Under the spot process $\mathcal{P}_\ell$ and the punishment scheme σ_0, the full compensation schedule $\bar{p}(R_{t-1}, \tilde{p}_{\ell t}|\eta + p)$ is:*

[24] These are bidders that become either primary or backup carriers.

[25] As one example, the Normal distribution has this strictly decreasing hazard rate property.

(i) strictly increasing in the carrier's rent $\eta + p$ for all $(R_{t-1}, \tilde{p}_{\ell t})$,

(ii) continuous in $\tilde{p}_{\ell t}$ for all R_{t-1} and $\eta + p$,

(iii) bounded below by $\eta + p + \kappa_\ell$ for all $(R_{t-1}, \tilde{p}_{\ell t})$.

Note that Assumption 2 is essentially an assumption on the underlying spot process and the shipper's punishment scheme. The most substantive assumption is 2(i), which can be numerically verified. Assumption 2(ii) holds under mild regularity conditions, and Assumption 2(iii) is satisfied if the shipper's punishment scheme σ_0 is strictly increasing in the carrier's rejection index R_{t-1}.

Theorem 1. *(Full identification) Under Assumptions 1 and 2, within the class of symmetric monotone equilibria, the model parameters $\{\sigma_0, \mathcal{P}_\ell, N, G_\ell^{\psi,\eta}, F_\ell, \kappa_\ell\}_\ell$ are identified.*

The spot process $\mathcal{P}_\ell$ is identified from the realized path of spot rates. The shipper's punishment scheme σ_0 is identified, since under Assumption 1, every Markov state $(R_{t-1}, \tilde{p}_{\ell t})$ is observed.[26] Given these instrumental objects $(\mathcal{P}_\ell, \sigma_0, N)$, we now identify the key model primitives: the distribution F_ℓ of operational costs, search cost κ_ℓ , and the joint distribution $G_\ell^{\psi,\eta}$ of the match-specific gains of shippers and carriers on each lane ℓ.

5.1 Carrier primitives from the repeated game

The key object in our identification of carrier primitives is carriers' tendency to accept a load at each Markov state of their relationships. Such tendency depends only on carriers' transformed rent $\eta_{ij\ell} + p_{ij\ell} + \kappa_\ell$ and the distribution $\tilde{F}_\ell$ of transformed costs $\tilde{c}_{j\ell t} = c_{j\ell t} + \kappa_\ell \sim \text{Normal}(\mu_\ell^c + \kappa_\ell, \sigma^c)$.[27] First, we argue that the level, shape, and variation of this acceptance tendency across carriers, as captured by their *acceptance schedules*, allow us to separately identify the distribution of their transformed rents and the distribution of their transformed costs. Then, we exploit the demand shifter z_ℓ to decompose transformed costs into operational and search costs.

Definition 2. (Acceptance schedules) Fix the shipper's strategy σ_0, the carrier's rent $\eta_{ij\ell} + p_{ij\ell}$, and lane characteristics $(\mathcal{P}_\ell, F_\ell, \kappa_\ell)$. Carrier j's acceptance schedule is that carrier's tendency to accept a load at each Markov state $(R_{t-1}, \tilde{p}_{\ell t})$,

$$\Pr(d_t = \text{accept} | R_{t-1}, \tilde{p}_{\ell t}) = \mathbf{1}\{\bar{p}(R_{t-1}, \tilde{p}_{\ell t}) \geq \tilde{p}_{\ell t}\} \tilde{F}_\ell(\bar{p}(R_{t-1}, \tilde{p}_{\ell t})).$$

[26] Specifically, in Assumption 1, (i) ensures that every level of spot rate is observed, regardless of the current rejection index; (ii) ensures that every Markov state is non-absorbing; and (iv) ensures a strictly positive probability of both acceptance and rejection in any Markov state.

[27] To see why, notice that the period-t payoff of carrier j can be rewritten as $(\eta_{ij\ell} + p_{ij\ell} + \kappa_\ell) - \tilde{c}_{j\ell t}$ if he accepts, as $\tilde{p}_{\ell t} - \tilde{c}_{j\ell t}$ if he services the spot market, and as 0 otherwise.

Figure 3: Optimal decisions and acceptance schedules for a fixed rejection index

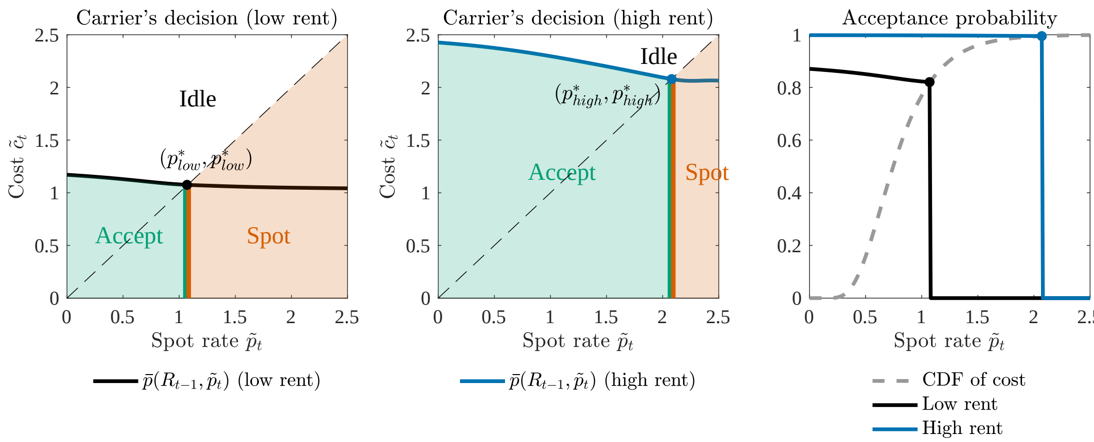

To build intuition as to why carriers' transformed costs and transformed rents can be separately identified from their acceptance schedules, consider a low-rent carrier and a high-rent carrier on the same lane, both with cost parameters $(\tilde{F}_\ell, \kappa_\ell)$. The first two panels of Figure 3 plot, for a fixed rejection index, the full compensations of the low-rent carrier (black line) and high-rent carrier (blue line) as well as their optimal decisions, in the space of spot rate (x axis) and transformed cost (y axis). As spot rate increases, the gap between the full compensation associated with each rent level and the spot rate shrinks, with the full compensation schedule crossing the 45-degree line at a critical point p^*. The carrier decides between *accept* and *idle* when spot rate is lower than p^*, between *spot* and *idle* when spot rate is higher than p^*, and is indifferent between *accept* and *spot* exactly at p^*. This means that the acceptance probability at p^*, which is the observed probability mass on the green vertical line, approximates the probability that the carrier chooses *spot* over *idle* locally to the right of this spot rate, which is the unobserved probability mass on the red vertical line,

$$\lim_{\tilde{p}_t \uparrow p^*} \Pr(d_t = \text{accept} | R_{t-1}, \tilde{p}_t) = \tilde{F}_\ell(p^*) = \lim_{\tilde{p}_t \downarrow p^*} \Pr(d_t = \text{spot} | R_{t-1}, \tilde{p}_t) > 0. \tag{7}$$

Moreover, since the carrier never accepts when $\tilde{p}_t > p^*$, the point $(p^*, \tilde{F}_\ell(p^*))$ manifests as a *jump point* identifiable from the carrier's acceptance schedule. The third panel of Figure 3 translates the optimal decisions of the low-rent and high-rent carriers in the first two panels into their acceptance schedules, whose distinct jump points trace the common distribution $\tilde{F}_\ell$ of transformed costs. As shown in the first and second panels, increasing a carrier's rent shifts out his full compensation. As a result, the critical points of the low-rent and high-rent carriers—and thus their corresponding jump points—are strictly ordered in their rents,

$p^*_{low} < p^*_{high}$.[28] In summary, the variation in carrier rent on the same lane identifies the common distribution $\tilde{F}_\ell$ of transformed costs on that lane and, conditional on the identified $\tilde{F}_\ell$, a carrier's acceptance schedule identifies his rent level.

While a carrier's acceptance schedule identifies his primitives, we face the empirical challenge that these schedules are not observed for short-lived relationships.[29] To overcome this challenge, we employ a two-step approach. The first step takes advantage of long-lasting relationships to identify the distribution of transformed costs. Then, taking this cost distribution as given, the second step pools all relationships on the same lane to identify the distribution of these carriers' rents from the pooled acceptance schedule. The intuition for the second step is as follows. Since the acceptance schedules associated with different levels of carrier rent jump to zero at different levels of spot rate, these schedules are linearly independent. Such linear independence implies that the pooled acceptance schedule identifies the mixture of acceptance schedules, which, in turn, identifies the distribution of carrier rent by the monotonicty of carriers' acceptance in their rents.[30] The following lemmas formalize our two-step approach.

Lemma 1. *(Identification of the distribution of transformed costs) A carrier's acceptance schedule on lane ℓ identifies at least one point on the distribution $\tilde{F}_\ell$ of transformed costs, and the variation in carrier rent across relationships on the same lane identifies $\tilde{F}_\ell$.*

Lemma 2. *(Identification of the distribution of transformed rents) Suppose that the instrumental objects and the distribution $\tilde{F}_\ell$ of transformed costs are identified. The distribution $[G_\ell^{\eta+p+\kappa}]^{1:N}$ of the transformed rents of winning carriers is identified.*

The final step in identifying carrier primitives is to decompose search costs from operational costs and pin down the potential causal link between search costs and spot market thickness. We do so by exploiting a demand shifter for spot volumes that is independent of unobserved cost factors across lanes. Subtracting the identified search costs from the distribution of transformed rents yields the distribution $[G_\ell^{\eta+p}]^{1:N}$ of winning carriers' rents.

[28] Our argument allows for the case that the full compensation equals the spot rate at multiple levels of spot rates. Appendix A provides the formal definition of *jump points* and conditions for their existence.

[29] On-path terminations of relationships occur due to demotion and auction events, resulting in short-lived relationships. For long-lasting relationships, we can identify the carrier's acceptance tendency conditional on the relationship surviving to the next period. Appendix E.1 provides details on how we correct for this selection.

[30] Kasahara and Shimotsu (2009) show, in general dynamic discrete choice models with Markov states, that linear independence of response functions is sufficient for identification of finite mixtures. Since we assume continuum support for the distribution of carrier rent, we present a direct proof of Lemma 2 in Appendix A.

Lemma 3. *(Cost decomposition) If $\tilde{F}_\ell$ is identified, then the variation in the demand shifter z_ℓ and in the average spot volume $Volume_\ell^{spot}$ identifies search costs $\kappa_\ell = \kappa(Volume_\ell^{spot})$ and the distribution F_ℓ of operational costs, up to a constant.*[31]

5.2 Shipper primitives from the auction

This section identifies the joint distribution of match-specific gains for shippers and carriers in realized relationships. Crucially, we exploit the existence of a strictly monotone mapping $\mathbf{b}_r : \eta + p \mapsto \psi - p$ between carriers' rents and shippers' rents in symmetric monotone equilibria. Adopting a strategy similar to Guerre et al. (2000), we pin down the derivative of this mapping from the first-order condition of carriers' bidding, albeit in the space of carrier rents rather than the space of observed bids. Another difference between our approach and theirs is that the initial condition for our mapping $\mathbf{b}_r$ is derived from the individual rationality constraint of the shipper (i.e., the auctioneer) rather than that of the carriers (i.e., the bidders). The reason for this is that long-term contracts in our setting allow the carriers to accept or reject any requests; carriers therefore always weakly benefit from relationships.

Finally, $\mathbf{b}_r$ maps the distribution of carrier rent, which was previously identified, to the distribution of shipper rent. Then, conditioning the joint distribution of shipper and carrier rents on contract rates pins down the joint distribution of match-specific gains. The following lemmas formalize our identification argument.

Lemma 4. *(Identification of the distribution of shipper rent) Suppose that a symmetric monotone equilibrium posits a one-to-one mapping $\boldsymbol{b}_r : \eta + p \mapsto \psi - \eta$ from carrier rent to shipper rent. Furthermore, suppose that the instrumental objects, the cost parameters (F_ℓ, κ_ℓ), and the distribution $[G_\ell^{\eta+p}]^{1:N}$ of winning carriers' rents are identified. Then the mapping $\boldsymbol{b}_r$ and, thus, the distribution $[G_\ell^{\psi-p}]^{1:N}$ of shippers' rents are identified.*

Lemma 5. *(Identification of the joint distribution of match-specific gains) Suppose that the instrumental objects, the cost parameters (F_ℓ, κ_ℓ), and the distribution $[G_\ell^{\eta+p|p}]^{1:N}$ of winning carriers' rents conditional on contract rates are identified. Then the joint distribution $[G_\ell^{\psi,\eta}]^{1:N}$ of shippers' and carriers' match-specific gains are identified.*

[31]While search costs are identified only up to a constant, our welfare conclusion relies mainly on how search costs scale (inversely) with the thickness of the spot market.

6 Estimation

In this section, we outline our estimation procedure, which closely follows the identification steps, and present our estimates of model primitives. These estimates suggest that long-term relationships generate large gains to the participating parties but exert substantial negative externalities on the spot market. We find that the total match quality of the median relationship is \$1.50/mile, of which \$0.85/mile comes from shipper and carrier match-specific gains and \$0.65/mile comes from savings on search costs for spot loads. Doubling spot market volume would reduce these search costs by half.

Following the logic of the identification argument outlined in the previous section, we use the combined relationship and spot data to sequentially recover (i) instrumental objects, (ii) cost parameters and spot markets' scale efficiency, and (iii) match-specific gains. For ease of interpretation, contract and spot rates, operational and search costs, and match-specific gains are estimated on a per-mile basis. Confidence intervals for our estimates are constructed from 50 full-estimation bootstraps, with sampling at the auction level. Appendix E explains our estimation procedure in detail.

6.1 Instrumental objects

Two instrumental objects are key to relationship dynamics: (i) the spot process and (ii) the shipper's punishment scheme. Since we do not observe the spot rate each carrier faces but only the time-specific mean $\tilde{p}_{\ell t}$ and standard deviation σ^{ζ}_{ℓ}, we assume that at the time of decision-making, carrier j faces spot rate $\tilde{p}_{\ell t} + \zeta_{j\ell t}$, with $\zeta_{j\ell t} \sim \text{Normal}(0, \sigma^{\zeta}_{\ell})$. We estimate an AR(1) process for the calendar-based spot rate ($\tilde{p}_{\ell\tau}$) normalized by the average spot rate,

$$\frac{\tilde{p}_{\ell\tau}}{\text{Rate}_{\ell}} = \rho_0 + \rho_1 \frac{\tilde{p}_{\ell\tau-1}}{\text{Rate}_{\ell}} + \epsilon_{\ell\tau}.$$

We then scale the estimated calendar-based process by the frequency of shipper-carrier interactions in a relationship to obtain the load-based spot ($\tilde{p}_{\ell t}$) process in that relationship.

For the shippers' punishment scheme, we jointly estimate a Probit specification,

$$\sigma_0(R_{t-1}, \tilde{p}_{\ell t}) = \Phi(\alpha_1 + \alpha_2 R_{t-1} + \alpha_3 \mathbf{X}_{ia\ell} + \alpha_4 R_{t-1}\mathbf{X}_{ia\ell} + \alpha_5(\tilde{p}_{\ell t} - \tilde{p}_{\ell 0})), \tag{8}$$

the initial rejection index R_0, and the daily decay parameter α on carriers' past rejections. Here, $\mathbf{X}_{ia\ell}$ includes the (log) frequency of shipper-carrier interactions and the coefficient of variation of the shipper's weekly volume; and $(\tilde{p}_{\ell t} - \tilde{p}_{\ell 0})$ captures changes in market

conditions since the start of the relationship.[32] Since the rejection index R_{t-1} might be correlated with unobserved variation in shippers' punishment scheme, we instrument for R_{t-1} with an analogously constructed index of past spot rates. Consistent with Harris and Nguyen (2024), we find that shippers punish carriers' rejections by increasing demotion probability in future periods; this punishment scheme is soft but generates meaningful economic incentives. Appendix F.2 presents our estimates of shippers' punishment scheme and other instrumental objects, including the discount factor and number of bidders.

6.2 Cost parameters and spot markets' scale efficiency

Parametric assumptions. Assume that transformed costs $\tilde{c}_{j\ell t} \sim \text{Normal}(\tilde{\mu}^c_{ia\ell}, \sigma^c)$ with auction-specific means $\tilde{\mu}^c_{ia\ell}$ and common variance σ^c. Moreover, assume that the mean $\tilde{\mu}^c_{ia\ell}$ of transformed costs can be decomposed into search cost κ_ℓ and the mean $\mu^c_{ia\ell}$ of operational costs as follows:

$$
\begin{aligned}
\tilde{\mu}^c_{ia\ell} = & \overbrace{\gamma_0 + \gamma_1 \ln(\text{Volume}^{\text{spot}}_\ell) + \gamma_2 \ln(\text{Volume}^{\text{spot}}_\ell) \ln(\text{Distance}_\ell)}^{\text{search cost } \kappa_\ell} \\
& + \underbrace{\gamma_3 \ln(\text{Distance}_\ell) + \gamma_4 \mathbf{X}_{ia\ell} + \gamma_5 \text{Tight}_a + \gamma_6 \text{Imbalance}_\ell + \nu^c_\ell + \epsilon^c_{ia\ell}}_{\text{mean operational cost } \mu^c_{ia\ell}} .
\end{aligned} \tag{9}
$$

In the specification of search costs, $\text{Volume}^{\text{spot}}_\ell$ is the average spot volume and Distance_ℓ is the average distance of lane ℓ. That is, we allow search costs to vary with spot market thickness, with this effect potentially varying with lane distance. In the specification of operational costs, Tight_a is an indicator of whether auction a occurs in a tight market phase, and Imbalance_ℓ captures the likelihood of a carrier finding a backhaul after completing the forehaul on lane ℓ. The former variable helps control for changes in carriers' opportunity costs across market phases; the latter mitigates concerns about spatial equilibrium effects.[33]

[32] As argued in Harris and Nguyen (2024), the frequency and timing consistency of load offers affect the value of a relationship, thus potentially affecting the harshness of the shipper's punishment scheme.

[33] In principle, search frictions might vary with market thickness (e.g., Vreugdenhil, 2023). Our specification subsumes such effect under the reduced-form term Tight_a, since our instrument has only cross-sectional variation and not temporal variation. We define by $\text{Imbalance}_\ell \equiv \ln(\text{Volume}^{\text{spot}}_{-\ell}) - \ln(\text{Volume}^{\text{spot}}_\ell)$ the volume imbalance on lane ℓ. Patterns of trade might affect the equilibrium movements of trucks, thus potentially correlating with unobserved cost shifters. For example, a thick lane might appear desirable not because search costs are lower on this lane but because it is connected to other thick lanes—this connectivity would help carriers reduce empty miles. In our setting, including Imbalance_ℓ in Equation (9) helps mitigate such spatial-equlibrium concern in a meaningful way because the ability to find backhauls is a key concern for carriers on long trips of at least 250 miles. Our estimates of cost parameters are robust to the inclusion of this imbalance measure.

Table 3: Estimates of cost determinants

Parameter	Variable	Estimate	95% bootstrap CI
γ_1	$\ln(\text{Volume}^{\text{spot}})$	-0.514	$(-0.756, -0.293)$
γ_2	$\ln(\text{Volume}^{\text{spot}}) \times \ln(\text{Distance})$	0.068	$(-0.186, 0.278)$
γ_3	$\ln(\text{Distance})$	-1.098	$(-1.440, -0.794)$
$\gamma_{4,1}$	Frequency	0.244	$(0.132, 0.359)$
$\gamma_{4,2}$	Inconsistency	0.660	$(-0.182, 1.307)$
γ_5	Tight	0.411	$(0.303, 0.573)$
γ_6	Imbalance	-0.244	$(-0.393, -0.117)$
σ^c	Cost variance	0.8	$(0.7, 1.1)$

Estimation procedure. We estimate cost parameters in three steps. First, we estimate transformed costs and rents by maximizing the likelihood of carriers' accept/reject decisions in relationships with at least 30 requests.[34] For each of these relationships and each set of parameter values $(\eta_{ij\ell} + p_{ij\ell} + \kappa_\ell, \tilde{\mu}^c_{ia\ell}, \sigma^c)$, we use a fixed-point algorithm to solve for the carrier's value function and optimal strategy (Rust, 1994). Second, using the estimates of $(\tilde{\mu}^c_{ia\ell})_{ia\ell:T_{ia\ell}\geq 30}$ from the previous step, we estimate Equation (9) by two-stage least squares using the predicted trade flows across states from Caliendo et al. (2018) as the instrument for spot volume. To complete the cost decomposition, we calibrate γ_0 to match the median operational cost to the industry estimate in Williams and Murray (2020).[35] Finally, we use observed relationship and lane characteristics to extrapolate $(\tilde{\mu}^c_{ia\ell})_{ia\ell:T_{ia\ell}\geq 30}$ to all lanes. Table 3 presents our estimates of cost parameters, and Figure 4 plots the estimated search costs and mean of operational costs across all relationships.

Large scale efficiency of the spot market. The estimates in Table 3 show that increasing the thickness of the spot market substantially reduces search costs. This supports the idea hypothesized by Kranton (1996) that the formation of long-term relationships can crowd out the spot market by making it thinner and less efficient. Our estimate of the scale efficiency parameter γ_1 is negative and economically significant, with no detectable effect of distance on scale efficiency. Specifically, we estimate that doubling the spot volume would decrease search costs on a lane by \$0.36/mile.

[34] Long-lasting relationships allow us to recover relationship-specific parameters but require correcting for selection on survival. Appendix E describes how we perform this correction on the likelihood function.

[35] Using an accounting approach, Williams and Murray (2020) estimate the marginal cost of trucking service to be \$1.55/mile, including fuel costs. Since long-term contracts typically separate payment on fuel costs as fuel surcharge, and spot rates in our data subtract fuel surcharge, we also subtract fuel surcharge (\$0.33/mile) from the accounting estimate. Note that this specific cost decomposition will not affect our welfare conclusion.

Figure 4: Volume-weighted distribution of search and operational costs

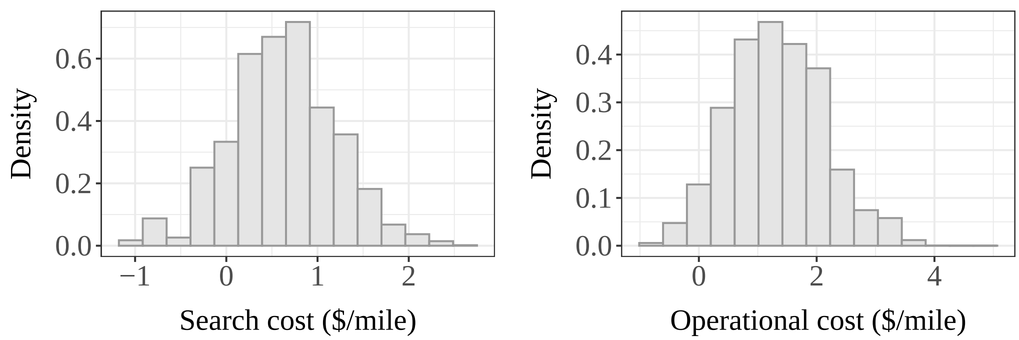

Determinants of operational costs. Table 3 also shows that per-mile operational costs are higher on shorter lanes and in relationships with higher volumes. Moreover, operational costs increase by \$0.41/mile in periods of sustained high demand and decrease by \$0.17/mile on lanes that are twice as easy to find backhauls. These factors contribute to the large variation in the means of operational costs across relationships in Figure 4.

6.3 Relationships' match-specific gains

Estimation procedure. Given our cost estimates, we recover match-specific gains in three steps. First, we estimate the distribution $G^{\eta+p}(\cdot|\tilde{\mathbf{X}}_{ia\ell})$ of carriers' rents conditional on an extended set of observable characteristics, $\tilde{\mathbf{X}}_{ia\ell} = \mathbf{X}_{ia\ell} \cup \{\text{Rate}_\ell, \text{Volume}_\ell^{\text{spot}}, \text{Distance}_\ell, \text{Tight}_a, \text{Imbalance}_\ell\}$. We use an EM algorithm (Train, 2008) to estimate this conditional distribution as a mixture of conditional Normal distributions. Second, we estimate the conditional mapping $\mathbf{b}_r|\tilde{\mathbf{X}}_{ia\ell} : \eta + p \mapsto \psi - p$ from carrier rents to shipper rents. Analogous to Guerre et al. (2000), we pin down the derivative of $\mathbf{b}_r$ from the first-order conditions of carriers' bidding in the space of carrier rents. For the initial condition of each mapping, we take the fifth percentile of the estimated distribution of carrier rent as the lowest carrier rent $\underline{r}_{ia\ell}$. The lowest shipper rent ensures that relationships have nonnegative match quality. That is, we set $\mathbf{b}_r(\underline{r}_{ia\ell}) = -\underline{r}_{ia\ell} - \kappa_\ell$.[36] Finally, conditioning rents on contract rates allows us to recover the joint distribution of match-specific gains.

Large and heterogeneous match-specific gains. We find large and heterogeneous match quality in realized relationships, with the shipper's match-specific gain accounting for a large fraction of total match quality.

[36]Since we only exploit individual rationality constraints, our estimates of shipper rents could be considered lower bounds. An approach that tries to rationalize the fact that shippers do not deviate to the spot market within the relationship would estimate even larger shipper rents.

Figure 5: Volume-weighted distribution of match-specific gains and match-quality

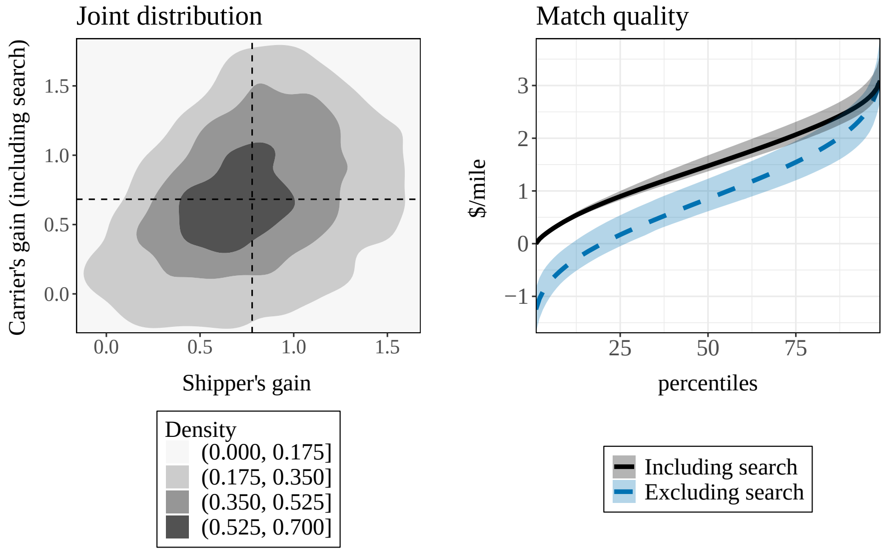

Note: The bands on the right panel are 95% bootstrap confidence intervals.

The left panel of Figure 5 plots the density of the joint distribution of carriers' match-specific gains including savings on search costs ($\eta + \kappa$) and shippers' match-specific gains (ψ) in realized relationships. Darker colors represent values of shippers and carriers' match-specific gains with higher density. We find that for the median relationship, the shipper's match-specific gain is \$0.78/mile and the carrier's gain is \$0.68/mile. However, only \$0.07/mile of this carrier's gain from a relationship transaction is due to the carrier's intrinsic gain; the rest comes from his savings on search costs.

The right panel of Figure 5 shows quantile plots for the distribution of match quality, including and excluding savings on search costs. Even when savings on search costs are excluded, 79% of relationships have positive intrinsic gains. This suggests that intrinsic relationship benefits play a more important role in driving the dominance of relationships in the current institution than does the thinness of the current spot market. However, these benefits are very heterogeneous across relationships, with intrinsic match quality increasing from \$0.15/mile to \$1.58/mile from the 25th percentile to the 75th percentile. Such large heterogeneity in match quality reflects the importance of auctions in facilitating the formation of high-value matches in the truckload setting.

7 Welfare analysis

In this section, we evaluate the performance of the current institution against alternatives. We find that the current relationships with fixed-rate contracts leave surplus on the table,

realizing only 64% of the relationship-level first-best surplus. However, solving this contractual friction without correcting for the market-thickness externalities of relationships on the spot market would not lead to welfare gains. In fact, during periods of high demand, this contractual friction within relationships helps balance relationship and spot performance. As a result, the current institution captures 92% of the market-level first-best surplus.

7.1 Conceptual framework and welfare calculation

Conceptual framework. Our welfare analysis focuses on two frictions in the US truckload setting: (i) the contractual friction within long-term relationships and (ii) the search friction in the spot market. The first friction could be solved by replacing the current fixed-rate contracts with *index-priced* contracts that transfer all of the rents from a relationship transaction to the carrier and select the carrier that bids the highest fixed fee.[37] Specifically, carriers would bid on $b^0_{ij\ell}$ for the following payment schedule,

$$p_{ij\ell}(\tilde{p}_{\ell t}) = \begin{cases} -\, b^0_{ij\ell} + \overbrace{\psi_{ij\ell} + \tilde{p}_{\ell t}}^{\text{“incentive”}} & \text{if carrier } j \text{ accepts} \\ \underbrace{-b^0_{ij\ell}}_{\text{“screening”}} & \text{if carrier } j \text{ rejects.} \end{cases}$$

While such index-priced contracts are optimal at the relationship level, market-level adoption of these contracts would result in even thinner and less efficient spot markets. If spot markets' scale efficiency is sufficiently large, the other extreme with no relationships but only spot transactions (*spot-only*) might be socially optimal.

More generally, the tradeoff between relationships' benefits and spot markets' scale efficiency can be assessed by tracing the highest welfare that can be achieved at each level of spot market's share in total transacted volume. Each such constrained first-best welfare can be implemented by combining the relationship-optimal index-priced contracts with an appropriate tax on relationship transactions. We will compare the current institution to the constrained first-best curve and zoom in on two counterfactuals: (i) *spot-only* (infinite tax) and (ii) *index-priced* (zero tax).

[37] Under a relationship-optimal index-priced contract, the carrier accepts when $c_{j\ell t} \leq (\psi_{ij\ell} + \eta_{ij\ell} + \tilde{p}_{\ell t})$ and remains idle otherwise. This acceptance rule delivers the relationship-level first-best because: (i) the shipper selects the relationship with the highest nonnegative match quality; (ii) in a relationship with nonnegative match quality, the carrier should never reject to service the spot market; and (iii) accepting at cost $c_{j\ell t}$ would realize the relationship's match quality $(\psi_{ij\ell} + \eta_{ij\ell})$ and save the shipper from paying $\tilde{p}_{\ell t}$ to a spot carrier.

Figure 6: Volume-weighted expected surplus from long-term relationships

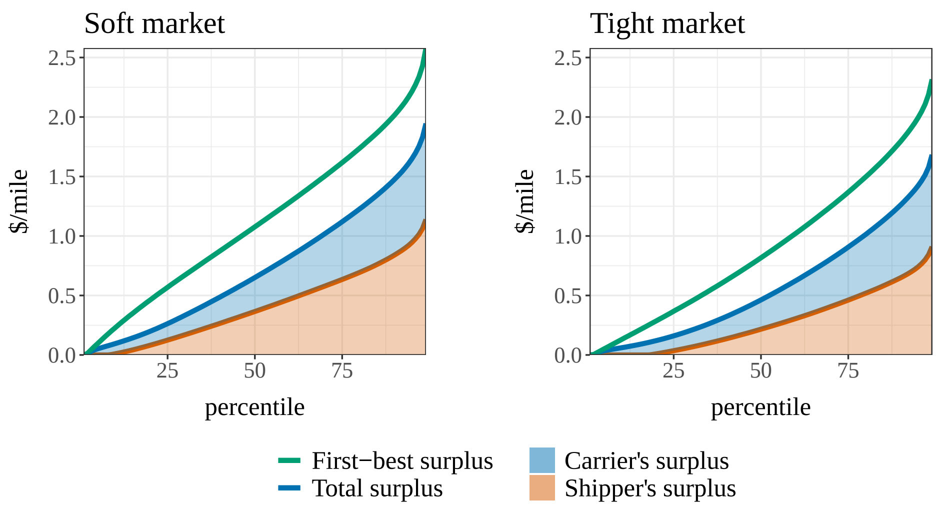

Welfare calculation. Since our relationship data covers only a subset of shippers, we cannot reliably estimate aggregate demand within relationships at the KMA-to-KMA level. Instead, we aggregate lanes into seven clusters. Within each cluster, we consider two full years, with the year 2016 representing a soft market and 2018 a tight market. Within each lane cluster and year, we calculate the current and counterfactual welfare in two steps. First, we recover the underlying aggregate demand and capacity, including: L relationships, all formed in $t = 0$; spot capacity C; and a vector $(D_t)_{t=1}^{52}$ of weekly direct spot demand. Second, we compute the welfare effects of different institutions while keeping fixed the underlying demand and supply factors $(L, D_t, C)_{t=1}^{52}$ and the distribution of match-specific gains and operational costs. What changes across market institutions are the equilibrium relationship volume, spot volume, spot rates, and search costs.[38] Thus, our welfare analysis answers the following question: fixing aggregate demand and capacity, how should transactions be allocated between relationships and the spot market?

7.2 Relationship-level welfare

Figure 6 plots the volume-weighted expected surplus from the current relationships (blue line) over spot transactions, how this surplus is split between shippers (red area) and carriers (blue area), against the relationship-level first-best surplus (green line).

Overall, we find that the current relationships generate a significant surplus of \$0.64/mile over spot transactions, which is shared equally between shippers and carriers. Moreover, the current relationships realize 64% the relationship-level first-best surplus, faring reasonably against *index-priced* contracts. This performance is similar across soft and tight markets

[38]See Appendix E.2 for the details of our welfare calculation.

Figure 7: Welfare change relative to the current institution

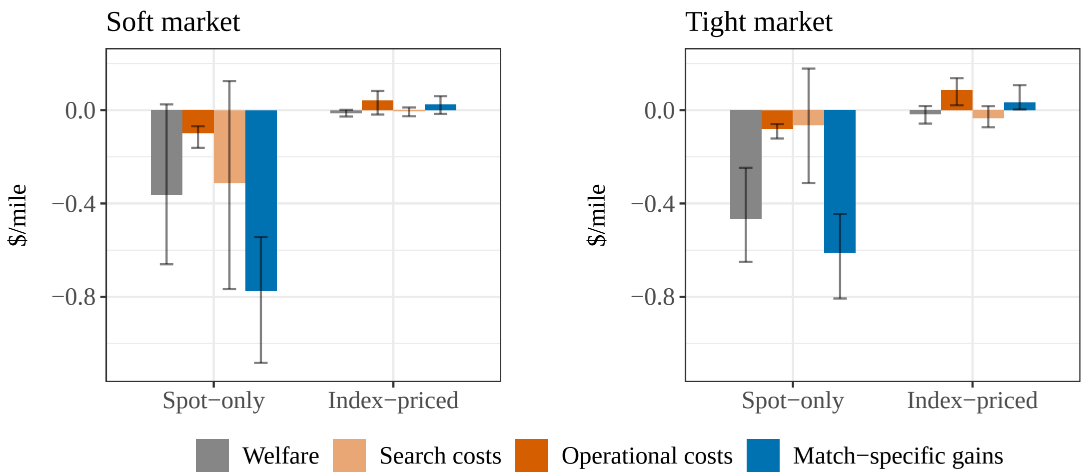

Note: 95% confidence intervals are constructed from 50 full-estimation bootstraps with auction-level sampling.

but varies greatly across relationships of different match quality.Specifically, a relationship at the 25th percentile of match quality achieves 45% of its potential relationship surplus as compared to 68% achieved by a relationship at the 75th percentile of match quality. That is, under the current contracts, spot markets crowds out low-value relationships in particular.

7.3 Market-level welfare

Figure 7 presents welfare changes ($/mile) for the *spot-only* and *index-priced* counterfactuals relative to the current institution, in which relationships employ fixed-rate contracts. To explore the sources of welfare gains and losses, we break down the welfare changes into three components: realized match-specific gains, operational costs, and search costs.

Spot-only counterfactual. We find that eliminating relationships would result in a substantial welfare loss relative to the current institution. This is because the reduction in search costs and operational costs, while large, is insufficient to compensate for the complete loss of match-specific gains from relationships. We estimate the welfare loss from this counterfactual to be $0.36/mile in a soft market and $0.47/mile in a tight market.

Index-priced counterfactual. We find similar market-level performance between the relationship-optimal index-priced contracts and the current fixed-rate contracts. This finding reflects the market-level tension between relationship and spot performance in our setting. While index-priced contracts optimize relationship performance, they result in thinner spot

markets, with search costs increasing by \$0.09/mile and \$0.22/mile in a soft and tight market, respectively. These higher search costs affect welfare through three channels. First, while searching is more costly, aggregate search costs are mostly unchanged, as the higher search costs are borne by a smaller number of carriers servicing the spot market. Second, these higher search costs widen the gap in realized operational costs between carriers servicing in relationships and those servicing in the spot market. This tends to hurt allocative cost efficiency. Third, stronger incentives to avoid search costs rationalize more relationship transactions with negative match quality. This dampens the gains from realizing more transactions with positive match quality.[39]

Market-level first-best. Our final exercise compares the current institution to the market-level first-best. The left panel of Figure 8 illustrates, for a cluster of lanes and the soft market phase, the constrained first-best curve and locates different institutions on this curve.[40] It indicates that both the *index-priced* counterfactual and the current institution with *fixed-rate* contracts have spot shares that are lower than the optimum, though the latter is closer to the first-best. This pattern holds for 12 out of our 14 pairs of lane clusters and market phases. Moreover, the portion of the constrained first-best curve between *index-priced* and the market-level first-best is relatively flat, suggesting that moving toward the first-best would require a large tax and deliver a relatively small welfare gain.[41]

The right panel of Figure 8 shows how far the current institution is from the market-level first-best across all clusters of lanes and market phases. Each point shows the current spot share against the first-best spot share, with lighter colors indicating that the current institution realizes a larger share of the market-level first-best surplus.[42] We find that the current institution realizes 88% of the market-level first-best surplus in a soft market and 95% in a tight market. This difference is due to the current institution approximating the first-best spot share much better in a tight market, when pressure from the spot market is high. That is, the incompleteness of the current contracts helps mitigate relationships' externalities on spot transactions, a mechanism that works well in periods of high demand.

[39] While the *index-priced* counterfactual would not imply significant changes in aggregate welfare, we show in Appendix F.3 that it would have substantial distributional consequences, hurting in particular spot carriers. In addition, fixed-rate contracts might provide insurance value that would be forgone under index-priced contracts; then, the relative performance of fixed-rate contracts would improve both at the relationship and market level.

[40] Technically, the current institution with fixed-rate contracts is not allocatively efficient. The fact that it resides so close to the constrained first-best curve suggests that the inefficiency of the current institution mostly comes from the suboptimal split between relationship and spot transactions.

[41] Across our 14 clusters, the first-best tax ranges from \$0.38/mile to \$0.54/mile.

[42] Since the *spot-only* counterfactual is unambiguously the worst performing institution, we use it as the baseline for normalization. Market-level surplus is then defined as the market-level welfare gain relative to this baseline.

Figure 8: Comparison to market-level first-best welfare

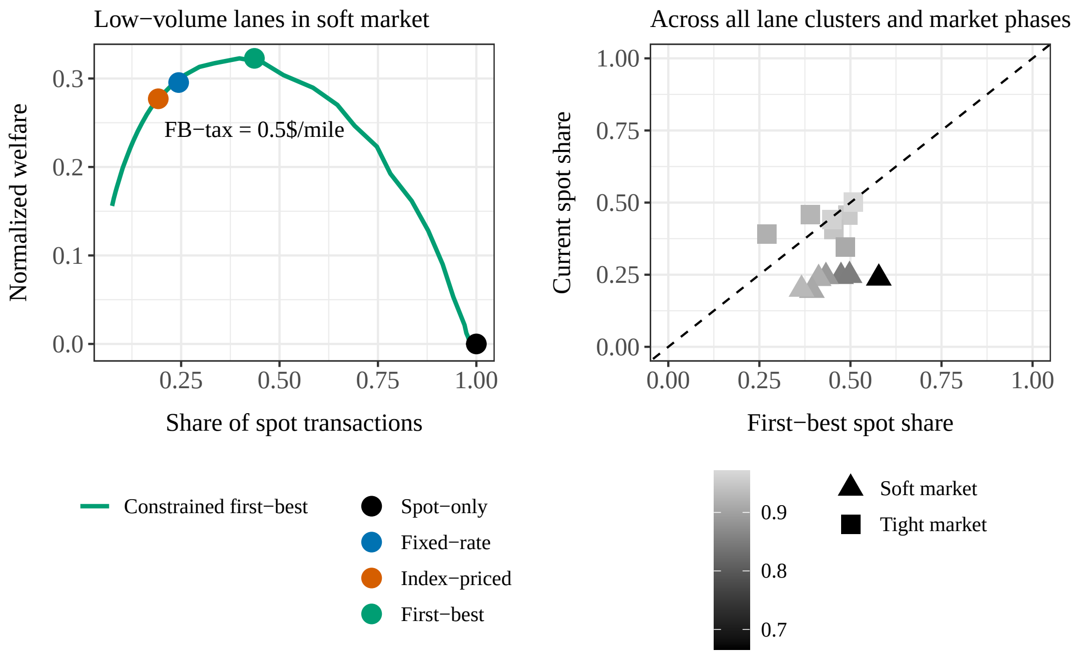

Notes: The color bar on the right panel shows the fraction of the market-level first-best surplus that is achieved by the current institution.

8 Conclusion

This paper studies the market-level tradeoffs between long-term relationships and the spot market. Using detailed data on the US for-hire truckload freight industry, we find that the two-way crowding-out effects between long-term relationships and the spot market, as hypothesized by Kranton (1996), play a crucial role in this setting. On the one hand, the current relationships have contracts that fix prices but not volumes, thus allowing spot temptation to crowd out relationships. On the other hand, spot markets have frictions that are worsened the more transactions are within relationships rather than in the spot market.

We build an empirical framework to identify the intrinsic gains from relationships and their externalities on the spot market, exploring how this tradeoff manifests under different market institutions. Our counterfactual analysis suggests that the benefits of relationships tend to outweigh their externalities. Moreover, the incompleteness of the current fixed-rate contracts helps mitigate such externalities, a mechanism that works particularly well in periods of high demand. Overall, the current institution achieves 92% of the market-level first-best surplus despite achieving only 64% of the relationship-level first-best surplus.

Our findings provide insights on potential market interventions. On the one hand, we find that a global shift from relationships to the spot market (e.g., by centralizing all transactions on a spot platform) could lead to substantial welfare loss. On the other hand, small

improvements in search technology (e.g., algorithmic search and pricing) could improve welfare, both directly by benefiting spot participants and indirectly by moving the equilibrium split between relationship and spot transactions toward the market-level first-best.

References

Acocella, Angela, Chris Caplice, and Yossi 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.

——— (2022a): "The end of set-it-and-forget-it pricing? Index-based freight contract design under uncertainty," MIT Center for Transportation & Logistics Working Paper Series.

——— (2022b): "When suppliers ignore their contracts: Quantifying the value of a customer in the buyer-supplier relationship," MIT Center for Transportation & Logistics Working Paper Series.

Adhvaryu, Achyuta, Vittorio Bassi, Anant Nyshadham, and Jorge A Tamayo (2020): "No line left behind: Assortative matching inside the firm," Working paper, National Bureau of Economic Research.

Admati, Anat R and Paul Pfleiderer (1988): "A theory of intraday patterns: Volume and price variability," *The Review of Financial Studies*, 1 (1), 3–40.

Allen, Jason and Milena Wittwer (2023): "Centralizing over-the-counter markets?" *Journal of Political Economy*, 131 (12), 3310–3351.

Baker, George, Robert Gibbons, and Kevin J Murphy (2002): "Relational contracts and the theory of the firm," *The Quarterly Journal of Economics*, 117 (1), 39–84.

Baker, George P and Thomas N Hubbard (2003): "Make versus buy in trucking: Asset ownership, job design, and information," *American Economic Review*, 93 (3), 551–572.

——— (2004): "Contractibility and asset ownership: On-board computers and governance in US trucking," *The Quarterly Journal of Economics*, 119 (4), 1443–1479.

Barron, Daniel, Robert Gibbons, Ricard Gil, and Kevin J Murphy (2020): "Relational adaptation under reel authority," *Management Science*, 66 (5), 1868–1889.

Blouin, Arthur and Rocco Macchiavello (2019): "Strategic default in the international coffee market," *The Quarterly Journal of Economics*, 134 (2), 895–951.

Bokher, Sam (2018): "Segments of U.S. trucking," Technical report.

Brancaccio, Giulia, Myrto Kalouptsidi, Theodore Papageorgiou, and Nicola Rosaia (2023): "Search frictions and efficiency in decentralized transport markets," *The Quarterly Journal of Economics*, 138 (4), 2451–2503.

Brugues, Felipe (2020): "Take the goods and run: Contracting frictions and market power in supply chains," Working paper, Brown University.

Buchholz, Nicholas (2022): "Spatial equilibrium, search frictions, and dynamic efficiency in the taxi industry," *The Review of Economic Studies*, 89 (2), 556–591.

Buchholz, Nicholas, Laura Doval, Jakub Kastl, Filip Matejka, and Tobias Salz (2024): "Personalized pricing and the value of time: Evidence from auctioned cab rides," Working paper.

Cajal-Grossi, Julia, Rocco Macchiavello, and Guillermo Noguera (2023): "Buyers' sourcing strategies and suppliers' markups in Bangladeshi garments," *The Quarterly Journal of Economics*, 138 (4), 2391–2450.

Caliendo, Lorenzo, Fernando Parro, Esteban Rossi-Hansberg, and Pierre-Daniel Sarte (2018): "The impact of regional and sectoral productivity changes on the US economy," *The Review of Economic Studies*, 85 (4), 2042–2096.

Caplice, Chris (2007): "Electronic markets for truckload transportation," *Production and Operations Management*, 16 (4), 423–436.

Casaburi, Lorenzo, Michael Kremer, and Ravindra Ramrattan (2019): "Crony capitalism, collective action, and ICT: Evidence from Kenyan contract farming," Working paper.

Castillo, Juan Camilo (2023): "Who benefits from surge pricing?" Working paper.

Frechette, Guillaume R, Alessandro Lizzeri, and Tobias Salz (2019): "Frictions in a competitive, regulated market: Evidence from taxis," *American Economic Review*, 109 (8), 2954–92.

Gaineddenova, Renata (2022): "Pricing and efficiency in a decentralized ride-hailing platform," Working paper.

Galenianos, Manolis and Alessandro Gavazza (2017): "A structural model of the retail market for illicit drugs," *American Economic Review*, 107 (3), 858–896.

Gavazza, Alessandro (2010): "Asset liquidity and financial contracts: Evidence from aircraft leases," *Journal of Financial Economics*, 95 (1), 62–84.

——— (2011): "Leasing and secondary markets: Theory and evidence from commercial aircraft," *Journal of Political Economy*, 119 (2), 325–377.

Gil, Ricard, Myongjin Kim, and Giorgio Zanarone (2021): "Relationships Under Stress: Relational Outsourcing in the US Airline Industry After the 2008 Financial Crisis," *Management Science*.

Guerre, Emmanuel, Isabelle Perrigne, and Quang Vuong (2000): "Optimal nonparametric estimation of first-price auctions," *Econometrica*, 68 (3), 525–574.

Harris, Adam and Thi Mai Anh Nguyen (2024): "Long-Term Relationships in the US Truckload Freight Industry," Working paper.

Holm, Seth (2020): "What is the difference between trucking contract and spot rates?" Technical report.

Hubbard, Thomas N (2001): "Contractual form and market thickness in trucking," *RAND Journal of Economics*, 369–386.

Kasahara, Hiroyuki and Katsumi Shimotsu (2009): "Nonparametric identification of finite mixture models of dynamic discrete choices," *Econometrica*, 77 (1), 135–175.

Kranton, Rachel E (1996): "Reciprocal exchange: A self-sustaining system," *American Economic Review*, 86 (4), 830–851.

Lagos, Ricardo (2003): "An analysis of the market for taxicab rides in New York City," *International Economic Review*, 44 (2), 423–434.

Macchiavello, Rocco and Ameet Morjaria (2015): "The value of relationships: Evidence from a supply shock to Kenyan rose exports," *American Economic Review*, 105 (9), 2911–45.

——— (2021): "Competition and relational contracts in the Rwanda coffee chain," *The Quarterly Journal of Economics*, 136 (2), 1089–1143.

Marcus, Alfred A (1987): "From market dominance to credible commitments: shippers' strategies in a deregulated trucking environment," *Transportation Journal*, 4–15.

Masten, Scott E (2009): "Long-term contracts and short-term commitment: Price determination for heterogeneous freight transactions," *American Law and Economics Review*, 11 (1), 79–111.

Nash, John (1953): "Two-person cooperative games," *Econometrica*, 128–140.

Nickerson, Jack A and Brian S Silverman (2003): "Why aren't all truck drivers owner-operators? Asset ownership and the employment relation in interstate for-hire trucking," *Journal of Economics & Management Strategy*, 12 (1), 91–118.

Ostria, Sergio J (2003): "Evaluation of US commercial motor carrier industry challenges and opportunities," Technical report.

Perrigne, Isabelle and Quang Vuong (2011): "Nonparametric identification of a contract model with adverse selection and moral hazard," *Econometrica*, 79 (5), 1499–1539.

Rosaia, Nicola (2023): "Competing platforms and transport equilibrium," Working paper.

Rose, Nancy L (1985): "The incidence of regulatory rents in the motor carrier industry," *The RAND Journal of Economics*, 299–318.

——— (1987): "Labor rent sharing and regulation: Evidence from the trucking industry," *Journal of Political Economy*, 95 (6), 1146–1178.

Rust, John (1994): "Structural estimation of Markov decision processes," *Handbook of econometrics*, 4, 3081–3143.

Startz, Meredith (2021): "The value of face-to-face: Search and contracting problems in Nigerian trade," PEDL Research Papers.

Tolvanen, Juha, Olivier Darmouni, and Simon Essig Aberg (2021): "Pulp Friction: The Value of Quantity Contracts in Decentralized Markets," Working paper.

Train, Kenneth E (2008): "EM algorithms for nonparametric estimation of mixing distributions," *Journal of Choice Modelling*, 1 (1), 40–69.

Vreugdenhil, Nicholas (2023): "Booms, busts, and mismatch in capital markets: Evidence from the offshore oil and gas industry," Working paper.

Williams, Nathan and Dan Murray (2020): "An analysis of the operational costs of trucking: 2020 update," Technical report.

Yang, Ron (2023): "(Don't) Take Me Home: Home Preference and the Effect of Self-Driving Trucks on Interstate Trade," Working paper.

Ying, John S (1990): "The inefficiency of regulating a competitive industry: productivity gains in trucking following reform," *The Review of Economics and Statistics*, 191–201.

Zahur, Nahim Bin (2022): "Long-Term Contracts and Efficiency in the Liquefied Natural Gas Industry," Working paper.

A Identification proof

Definition 3. (Jump points) Fix search cost κ and distribution F of operational costs. For carrier rent $\eta + p$ and rejection state R_{t-1}, define the jump point as the lowest spot rate above which the full compensation schedule is always lower than spot rate,

$$p^*(R_{t-1}|\eta + p) = \inf\{\hat{p} : \bar{p}(R_{t-1}, \tilde{p}_t|\eta + p) < \tilde{p}_t, \forall \tilde{p}_t > \hat{p}\}.$$

These jump points are identified from the carrier's acceptance schedule. Crucially, they are the levels of spot rate at which the carrier is indifferent between *accept* and *spot*. Two implications follow. First, the acceptance probability observed at the jump points proxies the carrier's decision to choose *spot* over *idle* were *accept* not to be an option; this hypothetical decision reveals information about the costs of servicing spot loads (Lemma 1). Second, the existence and ordering of jump points in carrier rent ensure linear independence of the acceptance schedules of carriers with different rents. This property is key to identifying the mixture of carrier rent pooling relationships of all duration (Lemma 2).

Lemma 6. *(Order of jump points) Fix search cost κ_ℓ, cost distribution F_ℓ, and rejection index R_{t-1}. Suppose that Assumptions 1 and 2 hold. Then the jump points are well-defined and $p^*(R_{t-1}|\eta + p) > p^*(R_{t-1}|\eta' + p')$ for every $\eta + p > \eta' + p' \geq 0$.*

Proof. To show that jump points are well-defined, we first show there exist $\hat{p}_{\text{low}}$ and $\hat{p}_{\text{high}}$ such that (i) $\bar{p}(R_{t-1}, \hat{p}_{\text{low}}) \geq \hat{p}_{\text{low}}$ and (ii) $\bar{p}(R_{t-1}, \tilde{p}_t) < \tilde{p}_t$ for all $\tilde{p}_t > \hat{p}_{high}$. The existence of $\hat{p}_{\text{low}}$ follows from Assumption 2(iii) that $\bar{p}(R_{t-1}, \tilde{p}_t)$ is bounded below. To show (ii), we show that $\bar{p}(R_{t-1}, \tilde{p}_t)$ is bounded above. Note that being in a relationship gives the carrier an additional option to accept a load and get a payoff of $\eta + p + \kappa_\ell - \tilde{c}_{\ell t}$, while not being in a relationship only gives the carrier the option to accept a spot load, which yields $\tilde{p}_{\ell t} - \tilde{c}_{\ell t}$, or to remain idle and get zero. Thus, the dynamic compensation of the carrier for an acceptance is bounded above by $\eta + p + \kappa_\ell$. It follows that $\bar{p}(R_{t-1}, \tilde{p}_{\ell t}) \leq \frac{\eta+p+\kappa}{1-\delta}$.

Suppose, for the sake of contradiction, that $p^*(R_{t-1}|\eta+p) \leq p^*(R_{t-1}|\eta'+p')$. Then, at $\tilde{p}_t = p^*(R_{t-1}|\eta'+p')$, we have $\bar{p}(R_{t-1},\tilde{p}_t|\eta'+p') = \tilde{p}_t \geq \bar{p}(R_{t-1},\tilde{p}_t|\eta+p)$, where the last inequality follows from the definition of p^*. However, under Assumption 2(i) that the full compensation in strictly increasing in carrier rent, $\bar{p}(R_{t-1},\tilde{p}_t|\eta'+p') < \bar{p}(R_{t-1},\tilde{p}_t|\eta+p)$. This yields a contradiction. $\square$

Proof of Lemma 1. (Identification of the distribution of transformed costs) By Lemma 6, that $\eta+p \neq \eta'+p'$ gives us at least two distinct jump points. Note that each jump point p^* represents a point on the distribution $\tilde{F} \sim \text{Normal}(\tilde{\mu}^c,\sigma^c)$ of transformed costs, where $\tilde{F}(p^*)$ is the observed acceptance probability at the jump point. Thus, two distinct jump points give us a system of linear equations

$$\frac{p^*(R_{t-1}|\eta+p)-\tilde{\mu}^c}{\sigma^c} = \Phi^{-1}\left(\Pr(d_t = \text{accept}|R_{t-1},\tilde{p}_t = p^*(R_{t-1}|\eta+p);\eta+p)\right)$$
$$\frac{p^*(R'_{t-1}|\eta'+p')-\tilde{\mu}^c}{\sigma^c} = \Phi^{-1}\left(\Pr(d_t = \text{accept}|R'_{t-1},\tilde{p}_t = p^*(R'_{t-1}|\eta'+p');\eta'+p')\right)$$

that exactly identify $(\tilde{\mu}^c,\sigma^c)$. $\square$

Proof of Lemma 2. (Identification of the distribution of carrier rent) Without loss of generality, we provide the proof for $\kappa = 0$. Fix a rejection index R_{t-1}. We exploit the following equality

$$\Pr(d_t = \text{accept}, R_{t-1},\tilde{p}_t) = \int \Pr(d_t = \text{accept}, R_{t-1},\tilde{p}_t|r = \eta+p)d[G^{\eta+p}]^{1:N}(r),$$

where the joint distribution of carriers' acceptance, rejection state and spot rate, both unconditional and conditional on the level of carrier rent, are either directly observed or identified. Furthermore, the acceptance probability conditional on carrier rent satisfies a key property that beyond its jump point, acceptance probability equals zero,

$$\Pr(d_t = \text{accept}, R_{t-1},\tilde{p}_t|r = \eta+p) = 0, \quad \text{for all } \tilde{p}_t > p^*(R_{t-1}|\eta+p).$$

Our task is to identify the mixture $[G^{\eta+p}]^{1:N}$.

Take any two distributions $[G^{\eta+p}]^{1:N}$ and $[\hat{G}^{\eta+p}]^{1:N}$ with absolutely continuous densities $[g^{\eta+p}]^{1:N}$ and $[\hat{g}^{\eta+p}]^{1:N}$ that are not everywhere the same. Let $\bar{r} = \inf\{r' : [g^{\eta+p}]^{1:N}(r) = [\hat{g}^{\eta+p}]^{1:N}(r), \forall r > r'\}$ and suppose that $[g^{\eta+p}]^{1:N}(\bar{r}^-) > [\hat{g}^{\eta+p}]^{1:N}(\bar{r}^-)$. The continuity of $[g^{\eta+p}]^{1:N}$ and $[\hat{g}^{\eta+p}]^{1:N}$ further implies that for some $\epsilon > 0$, $[g^{\eta+p}]^{1:N}(\bar{r}) > [\hat{g}^{\eta+p}]^{1:N}(\bar{r})$ for all

$r \in [\bar{r}-\epsilon, \bar{r}]$. It follows from Lemma 6 that

$$\int_{p^*(R_{t-1}|\eta+p=\bar{r}-\epsilon)}^{\infty} \Pr(d_t = \text{accept}, R_{t-1}, \tilde{p}_t > p^*(R_{t-1}|\eta+p=\bar{r}-\epsilon)|\eta+p=r) d[G^{\eta+p}]^{1:N}(r)$$
$$> \int_{p^*(R_{t-1}|\eta+p=\bar{r}-\epsilon)}^{\infty} \Pr(d_t = \text{accept}, R_{t-1}, \tilde{p}_t > p^*(R_{t-1}|\eta+p=\bar{r}-\epsilon)|\eta+p=r) d[\hat{G}^{\eta+p}]^{1:N}(r).$$

That is, two different distributions generate different acceptance probability on some range of spot rates, completing the proof that $[G^{\eta+p}]^{1:N}$ is nonparametrically identified. □

Proof of Lemma 4. (Identification of the distribution of shipper rent) In a symmetric monotone equilibrium with an effective bidding function $\mathbf{b}$ and a rent function $\mathbf{r}: \theta \mapsto r = \theta - \mathbf{b}(\theta)$, there exists a unique monotone mapping $\mathbf{b}_r : r \mapsto b$ defined by $\mathbf{b}_r(r) = \mathbf{b}(\mathbf{r}^{-1}(r))$. Thus, the first-order condition for the optimal bidding of carrier j with type $\theta \geq \underline{\theta}$ can be written in his equilibrium rent $r = \mathbf{r}(\theta) \geq \underline{r}$,

$$(N-1)\frac{g^{\eta+p}(r)}{G^{\eta+p}(r)} = \frac{\frac{\partial}{\partial r}V(R_0, \tilde{p}_0|r)}{V(R_0, \tilde{p}_0|r) - \mathbf{E}[\underline{V}(\tilde{p}_1)|\tilde{p}_0]}\mathbf{b}'_r(r). \tag{10}$$

Notice that the carrier's expected payoff as a function of carrier's rent r is identified from the incentive scheme σ_0, cost distribution F and search cost κ. To pin down $\mathbf{b}'_r$ from Equation (10), it remains to show that $G^{\eta+p}$ is identified on $[\underline{r}, \infty)$.

For any rent level $r > \underline{r}$, the distribution of winning carriers' rents satisfies

$$[G^{\eta+p}]^{1:N}(r) = \frac{[G^{\eta+p}(r)]^N - [G^{\eta+p}(\underline{r})]^N}{1 - [G^{\eta+p}(\underline{r})]^N}. \tag{11}$$

Note that $G^{\eta+p}(\underline{r})$ is identified from the distribution of the number of effective bidders, which is Binomial$(N, 1-G^{\eta+p}(\underline{r}))$. It follows that $G^{\eta+p}$ is identified on $[\underline{r}, \infty]$ from the distribution of winning carriers' rents in Equation (11).

For the initial condition of $\mathbf{b}_r$, we exploit the shipper's IR constraint,

$$U(R_0, \tilde{p}_0|\mathbf{b}(\underline{r}), \underline{r}) = \mathbf{E}[\underline{U}(\tilde{p}_1)|\tilde{p}_0]. \tag{12}$$

Note that $\underline{r}$ is identified as the lowest point in the support of $G_r^{\eta+p}$, and the RHS of Equation (12) is identified from the spot process. Thus, Equation (12), with the LHS being strictly increasing in the shipper's rent, pins down $\mathbf{b}_r(\underline{r})$. □

Online Appendix

B Additional institutional details

Table 4: Example load offers: Shipper Z, lane City X - City Y (on June 1, 2018)

Type	Order	Carrier	Rate	Decision
Primary carrier	1	A	1.60	Reject
Backup carriers	2	B	1.44	Reject
	3	C	1.72	Accept
	4	D	1.89	

Note: The first request was sent to carrier A at the contracted rate of \$1.60/mile; since A rejected, a request was sent to B at the contracted rate of \$1.44/mile; since B also rejected, a request was sent to C at \$1.72/mile; since C accepted, no request was sent to D.

Figure 9: Temporal variation of contract and spot rates

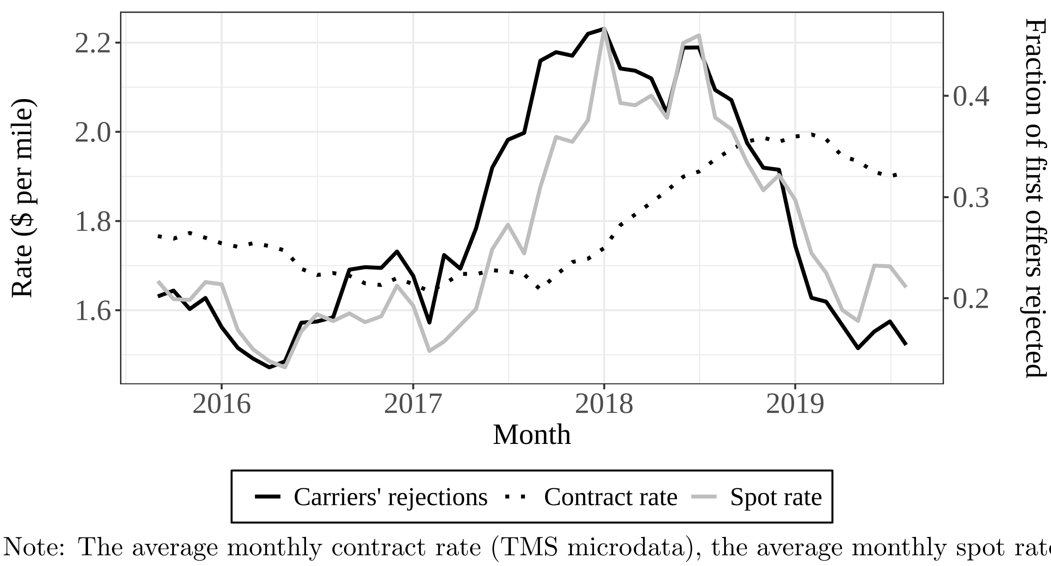

Note: The average monthly contract rate (TMS microdata), the average monthly spot rate (DAT data), and the rejection rate (TMS microdata) are all volume-weighted averages on the same set of lanes covered by the TMS microdata. The monthly rejection rate is the fraction of loads rejected by the primary carriers. Source: Harris and Nguyen (2024).

C Data Construction

This section describes the construction of key variables for the analysis of long-term relationships: per-mile contract rates, primary status, demotion events, and auction events. In addition to the observed routing guide for each load offer, we exploit a complementary data set that records the timestamps of shippers' input into the TMS. These timestamps provide the candidates for demotion and auction events. We refer to the period between two consecutive timestamps of a shipper on a lane as a "date-range."

Contract rates. Shippers seeking long-term relationships define lanes at geographical levels finer than KMA to KMA, sometimes as fine as warehouse-to-warehouse. Shippers can also bundle multiple origin-destination pairs in close proximity into a single "lane", using the same contract. This means that if a contract specifies a linehaul rate (total payment for a trip) on such a lane, the carrier's per-mile payment would vary with the distance of specific trips. On the other hand, if a contract specifies a per-mile rate, the carrier's total payment would vary with trips' distances. To match the unit of spot rates, we construct per-mile contract rates for specific trips and take the median of these rates within a date-range as the fixed contract rate.

Primary status. We infer carriers' status from the fact that primary carriers are generally the first to receive shippers' requests. Exceptions are typically due either to pre-specified capacity constraints that both the shipper and the carrier agreed to or to multiple primary carriers sharing the same lane. In such instances, we assign primary status to the carrier with the most first requests within a date-range.

Auction and demotion events. Our data do not include direct indicators of auction and demotion events. However, we can infer these events from observed changes in contract rates (for auction events) and the orderings of requests (for demotion events).

We treat the beginning of the current date-range as an auction event if we observe changes in the contract rates of the same carriers (either linehaul or per-mile) from what they were in the previous date-range. A secondary indicator of an auction event would be a completely new carrier replacing the previous primary carrier. However, there is a tradeoff in using this indicator. On the one hand, not using this indicator risks missing some auction events because carriers sometimes reuse their bids. On the other hand, using this indicator risks assigning an auction event to what is actually a demotion event, since some backup carriers may not appear in the routing guide. We judge that the second risk is less severe and therefore use both indicators to detect auction events.

Figure 10: Number of identified auction and demotion events

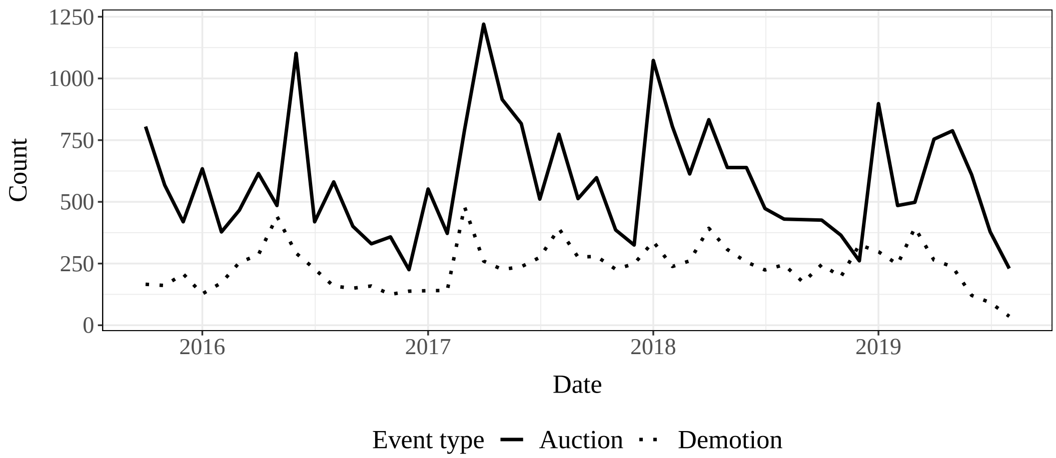

To identify demotion events, we identify instances in which the current primary carrier is replaced by another carrier within the same contract period (that is, between two auction events). Measurement errors in our constructed indicator of demotion events can come from measurement errors in our constructed primary status or in our indicators of auction events.

Figure 10 plots the number of identified auction and demotion events in each month-year in our sample period. Auctions appear to occur at random over time. There are spikes in the number of identified auction events, reflecting the fact that shippers tend to hold auctions on multiple lanes simultaneously. Identified demotion events are relatively evenly distributed over time and do not seem to correlate with identified auction events. This suggests that our data construction does a reasonably good job at separating the two types of events.

Carrier types. The trucking industry uses two different sets of codes as public identifiers for carriers: the Standard Carrier Alpha Code (SCAC), maintained by the National Motor Freight Traffic Association (NMFTA), and US DOT registration codes issued by the US Department of Transportation. We map the SCAC variable in our data set to US DOT codes using a conversion table from the NMFTA. We then map US DOT codes to carriers' registration at the Department of Transportation for the year 2020. This method matches 90% of carriers in our microdata. We divide the matched carriers into two types: (i) asset-owner carriers, defined as carriers with some asset ownership, and (ii) brokers, defined as carriers without any asset ownership. In Appendix F.1, we show that relationship performance varies substantially across these two groups of carriers.

Our main analysis uses a merged relationship and spot data set that restricts to asset-

Table 5: Summary statistics of the merged relationship and spot data

	Count	Variable	Mean	Standard deviation (SD)	
Relationships	12,586	Relationship duration (requests)	52	124	
		Frequency (requests/month)	11	18	
		Likelihood of eventual demotion	25%		
Auctions	10,262	Auction period (days)	381	250	
		Likelihood of ≥ 1 demotion	26%		
				Temporal	Spatial
KMA-KMA lanes	5,679	Contract rate (\$/mile)	1.83	0.12	0.50
		Spot rate (\$/mile)	1.87	0.23	0.50
		Spot volume (load posts/month)	1142	472	1136

owner carriers. Table 5 provides summary statistics for this data set.

D Simulations and other proofs

D.1 Monotonicity in rents

Assumption 3. *There exists $\underline{b} \in \mathbb{R}$ such that for all $b \geq \underline{b}$, $U(R_0, \tilde{p}_0|r, b) \geq \underline{U}(R_0, \tilde{p}_0)$ for all $r \geq 0$ and $U(R_0, \tilde{p}_0|r, b)$ is increasing in $r \geq 0$ and $b \geq \underline{b}$.*

The intuition for this assumption is that if the shipper's rent is sufficiently high, then fixing her rent, the shipper benefits from the carrier having a higher rent and thus accepting more frequently. While intuitive, this statement relies on the specifics of how the carrier's rent affects his path of play and how that path of play is correlated with the realized path of spot rates.

Figure 11 confirms that under our estimates of the spot process and shippers' incentive scheme, substantive assumptions on carriers' full compensation and shippers' expected payoffs are satisfied. The left panel shows that the carrier's full compensation is increasing in his rent (Assumption 2(i)). The right panel shows that the shipper's expected payoff is increasing in both her rent and the carrier's rent (Assumption 3).

D.2 Existence of a symmetric monotone equilibrium

Proposition 1. *Under Assumptions 1-3, a symmetric monotone equilibrium exists.*

Proof. We prove the existence of symmetric monotone equilibria in two steps. First, we construct a monotone equilibrium in an auxiliary game in which only match quality matters.

Figure 11: Monotonicity of carriers' full compensation and shippers' expected payoffs in rents

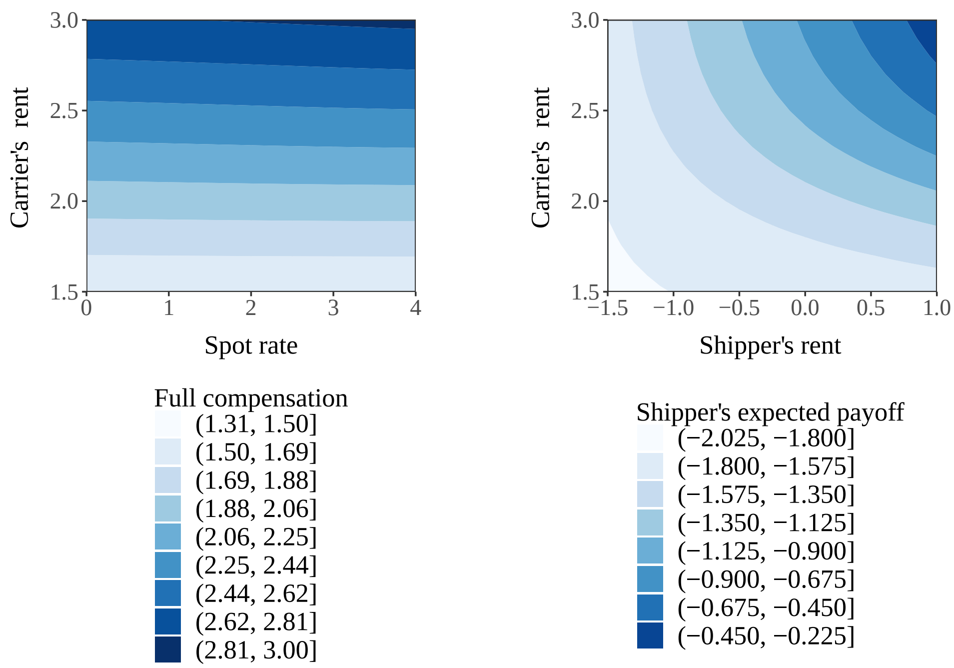

Second, we derive a symmetric monotone equilibrium in the original game from the monotone equilibrium of the auxiliary game.

Step 1: A monotone equilibrium of an auxiliary game.

Consider a bidding game where each carrier j has private information about his match-quality with the shipper, θ_{ij}. Each carrier submits a bid b_{ij} and the shipper chooses the carrier with the highest bid subject to reserve price $\underline{b}$. Here, $\underline{b}$ is the lowest level of shipper's rent that satisfies Assumption 3. The carrier that wins this auction gets expected payoff $V(R_0, \tilde{p}_0|\theta_{ij} - b_{ij})$.

In this game, there exists a strictly increasing bidding function $\mathbf{b} : \theta_{ij} \mapsto b_{ij}$ such that

$$\mathbf{b}(\theta_{ij}) = \arg\max_b [G^\theta(\mathbf{b}^{-1}(b))]^{N-1}(V(R_0, \tilde{p}_0|\theta_{ij} - b) - \mathbf{E}[\underline{V}(\tilde{p}_1)|\tilde{p}_0]).$$

Note that a relationship strictly benefits the carrier if and only if the carrier's rent is strictly positive. Thus, in this equilibrium, the lowest match quality of a winning carrier gives zero rent to that carrier, $\mathbf{b}(\underline{\theta}) = \underline{\theta}$. That is, individual rationality binds for the carrier with the lowest match quality. Moreover, a carrier with match quality $\theta > \underline{\theta}$ has a strictly positive rent, since he would otherwise strictly benefit from deviating to a lower bid. Thus, $\mathbf{r}'(\underline{\theta}) > 0$. Denote by $\mathbf{r} : \theta_{ij} \mapsto \theta_{ij} - \mathbf{b}(\theta_{ij})$ the function that maps the carrier's match quality to his

rent. We have $\mathbf{r}(\underline{\theta}) = 0$, and for all $\theta \geq \underline{\theta}$, $\mathbf{r}(\theta) > 0$ and $\mathbf{b}'(\theta) + \mathbf{r}'(\theta) = 1$. We want to show that $\mathbf{r}$ is strictly increasing.

The first-order condition of the carrier's bidding satisfies that for all $\theta > \underline{\theta}$,

$$(N-1)\frac{g^{\theta}(\theta)}{G^{\theta}(\theta)} = \frac{\frac{\partial}{\partial r}V(R_0, \tilde{p}_0 | r = \mathbf{r}(\theta))}{V(R_0, \tilde{p}_0 | r = \mathbf{r}(\theta)) - \mathbf{E}[\underline{V}(\tilde{p}_1)|\tilde{p}_0]}\mathbf{b}'(\theta).$$

Suppose that for some $\theta \geq \underline{\theta}$, $\mathbf{r}'(\theta) \leq 0$ and consider two cases: (i) there exists a strict interval on which $\mathbf{r}'(\theta) = 0$, and (ii) there is no such interval. In case (i), there exist $\theta_1 < \theta_2$ such that $\mathbf{r}(\theta_1) = \mathbf{r}(\theta_2)$ and $\mathbf{r}'(\theta_1) = \mathbf{r}'(\theta_2)$. In case (ii), there exist $\theta_1 < \theta_2$ such that $\mathbf{r}(\theta_1) = \mathbf{r}(\theta_2)$ and $\mathbf{r}'(\theta_1) > 0 > \mathbf{r}'(\theta_2)$. In either case, we have $0 < \mathbf{b}'(\theta_1) \leq \mathbf{b}'(\theta_2)$. Then under Assumption 1(iii) that G^{θ} has strictly decreasing hazard rate, we have

$$\begin{aligned}\frac{\frac{\partial}{\partial r}V(R_0, \tilde{p}_0 | r = \mathbf{r}(\theta_1))}{V(R_0, \tilde{p}_0 | r = \mathbf{r}(\theta_1)) - \mathbf{E}[\underline{V}(\tilde{p}_1)|\tilde{p}_0]} &= \frac{g^{\theta}(\theta_1)}{G^{\theta}(\theta_1)}[\mathbf{b}'(\theta_1)]^{-1} \\ &> \frac{g^{\theta}(\theta_2)}{G^{\theta}(\theta_2)}[\mathbf{b}'(\theta_2)]^{-1} = \frac{\frac{\partial}{\partial r}V(R_0, \tilde{p}_0 | r = \mathbf{r}(\theta_2))}{V(R_0, \tilde{p}_0 | r = \mathbf{r}(\theta_2)) - \mathbf{E}[\underline{V}(\tilde{p}_1)|\tilde{p}_0]}.\end{aligned}$$

This is a contradiction, completing the proof that $\mathbf{r}(\theta)$ is strictly increasing in θ.

Step 2: Symmetric monotone equilibrium.

We now map the monotone equilibrium of the auxiliary game to a symmetric monotone equilibrium of the original bidding game. Note that for a carrier j, if the shipper chooses the carrier with the highest effective bid (or proposed shipper's rent) and other carriers bid according to $\mathbf{b}$, then carrier j has no incentive to deviate from bidding according to $\mathbf{b}$. It remains to show the shipper's selection rule in the auxiliary game is optimal in the original bidding game.

Under Assumption 3 and by the choice of $\underline{b}$ in the auxiliary game, we have for all $\theta \geq \underline{\theta}$,

$$U(R_0, \tilde{p}_0 | \mathbf{b}(\theta), \mathbf{r}(\theta)) \geq \mathbf{E}[\underline{U}(\tilde{p}_1)|\tilde{p}_0]$$

and $U(R_0, \tilde{p}_0 | \mathbf{b}(\theta), \mathbf{r}(\theta))$ is increasing in θ. This means that by choosing the carrier j with the highest bid $b_{ij} \geq \mathbf{b}(\underline{\theta})$, the shipper maximizes her expected payoff from the relationship and never receives an expected payoff lower than her outside option of always going to the spot market. $\square$

E Estimation details

E.1 Estimation of model primitives

Step 1. Estimate shippers' strategy and instrumental objects

We parameterize the decay parameter in a relationship as a common daily decay parameter, α, scaled by the frequency of shipper-carrier interactions in that relationship. Parameters of the shippers' strategy include the daily decay parameter α, the initial rejection index R_0, and parameters $(\alpha_k)_{k=1}^5$ for the effects of the rejection index R_{t-1}, relationship characteristics $\mathbf{X}_{ia\ell}$, and market condition $(\tilde{p}_{\ell t} - \tilde{p}_{\ell 0})$ on demotion probability. We estimate these parameters in two steps. First, we estimate a linear counterpart of Equation (8), searching in an outer loop for the values of (α, R_0) that minimize the GMM objective. In the inner loop, we instrument for the rejection index R_{t-1} with (i) an index of past spot rates analogously constructed from the candidate values of (α, R_0) and (ii) the average spot rates in each of the last four weeks. The second set of instruments give us identification power to pin down (α, R_0). Given the estimates of (α, R_0) from the GMM procedure, we estimate $(\alpha_k)_{k=1}^5$ from the Probit specification using the same set of instruments. Table 7 presents our estimates of shippers' strategy in both the linear (GMM) and Probit (IV Probit) specifications.

We calibrate the daily discount rate to 0.992, under the assumption that (i) shippers and carriers are patient and (ii) auction periods end randomly with an estimated average duration of 381 days.

The final instrumental object is the distribution of the number of bidders, which captures the competitiveness of the auctions for long-term contracts. Our empirical model differs from the theoretical model in Section 4 in that we allow the number of bidders in an auction to be stochastic, $N_a \sim \text{Binomial}(N, q)$. Then the number of bidders who pass the shipper's individual rationality constraint and become either primary or backup carriers is $n_a \sim \text{Binomial}(N, \tilde{q})$, where $\tilde{q} = q(1 - G^{\eta+p}(\underline{r}))$. To facilitate estimation, we assume common parameters $(N, \tilde{q})$ for all auctions.

The empirical challenge is that we do not directly observe n_a; instead, we observe its lower bound, the number $\hat{n}_a$ of carriers that receive at least one request within the auction period. We thus estimate $(N, \tilde{q})$ via a calibration exercise. First, we take the maximum of $\hat{n}_a$ across all auctions as an estimate of N. Second, given this estimate of N and a value $\tilde{q}$, we calibrate the distribution of $\hat{n}_a$, taking into account the number of loads within an auction period, the probability that a load is rejected conditional on being rejected by higher-ranked carriers, and the current spot rate. Matching the mean of this simulated distribution to its empirical counterpart pins down $\tilde{q}$. We estimate that the number of effective bidders is distributed as $\text{Binomial}(20, 0.17)$, which has an average of 3.4 bidders per auction.

Step 2. Estimate carriers' transformed costs and rents

Carriers' transformed costs and rents are estimated by maximizing the likelihood of their observed accept/reject decisions. Given the observed standard deviation σ_ℓ^ζ of spot rate and a set of parameter values $(\tilde{\mu}_{ia\ell}^c, \sigma^c, \eta_{ij\ell} + p_{ij\ell} + \kappa_\ell)$, the likelihood contribution of relationship $ij\ell$ is

$$\begin{aligned}
&\mathcal{L}\left((\mathbf{1}\{d_{ij\ell t} = \text{accept}\}, R_{ij\ell t-1}, \tilde{p}_{\ell t})_{t=1}^{T_{ij\ell}}; \sigma_\ell^\zeta, \tilde{\mu}_{ia\ell}^c, \sigma^c, \eta_{ij\ell} + p_{ij\ell} + \kappa_\ell\right) \qquad (13)\\
&\propto \prod_{t=1}^{T_{ij\ell}} \prod_{D \in \{\{\text{accept}\},\{\text{idle,spot}\}\}} \Pr(d_{ij\ell t} \in D | R_{ij\ell t-1}, \tilde{p}_{\ell t}; \sigma_\ell^\zeta, \tilde{\mu}_{ia\ell}^c, \sigma^c, \eta_{ij\ell} + p_{ij\ell} + \kappa_\ell)^{\mathbf{1}\{d_{ij\ell t} \in D\}}.
\end{aligned}$$

Here, the acceptance probability in each period t of relationship $ij\ell$ is

$$\begin{aligned}
&\Pr(d_{ij\ell t} = \text{accept} | R_{ij\ell t-1}, \tilde{p}_{\ell t}; \sigma_\ell^\zeta, \tilde{\mu}_{ia\ell}^c, \sigma^c, \eta_{ij\ell} + p_{ij\ell} + \kappa_\ell) \qquad (14)\\
=&\Phi\left(\frac{\bar{p}(R_{ij\ell t-1}, \tilde{p}_t | \eta_{ij\ell} + p_{ij\ell} + \kappa_\ell; \sigma_\ell^\zeta, \tilde{\mu}_{ia\ell}^c, \sigma^c) - \tilde{p}_{\ell t}}{\sigma_\ell^\zeta}\right)\\
&\times \Phi\left(\frac{\bar{p}(R_{ij\ell t-1}, \tilde{p}_t | \eta_{ij\ell} + p_{ij\ell} + \kappa_\ell; \sigma_\ell^\zeta, \tilde{\mu}_{ia\ell}^c, \sigma^c) - \tilde{\mu}_{ia\ell}}{\sigma^c}\right),
\end{aligned}$$

A key input to the carrier's acceptance probability is his full compensation $\bar{p}$. For each set of parameter values, we solve for the carrier's full compensation by a fixed-point algorithm iterating between Equation (4) and the Bellman equation of the carrier's value function.

We exploit carriers' likelihood contribution in two separate steps: first, to estimate transformed costs using long-lasting relationships ($T \geq 30$) and second, to estimate the distribution of carrier rents pooling relationships of all lengths.

Step 2.1. Maximum-likelihood estimation of long-lasting relationships

In this first step, we restrict to relationships with at least 30 offers to estimate the parameters of transformed costs and rents $\{(\tilde{\mu}_{ia\ell}^c, \sigma^c, \eta_{ij\ell} + p_{ij\ell} + \kappa_\ell)\}_{T_{ij\ell \geq 30}}$ in these relationships. Focusing on long-lasting relationships have two potential selection issues: long-lasting relationships tend to have either (i) high rents or (ii) low cost draws. Since a carrier's rent in his relationship is a free parameter when we estimate cost parameters, selection due to (i) is not a concern. Selection due to (ii) can be corrected by conditioning the likelihood of observed

carriers' decisions on whether these carriers survive to the next period,

$$
\begin{aligned}
&\Pr(d_{ij\ell t} = \text{accept}|R_{ij\ell t-1}, \tilde{p}_{\ell t}, \text{surviving}) \\
=&\frac{[1-\sigma_0(\alpha R_{ij\ell t-1}, \tilde{p}_t)] \Pr(d_{ij\ell t} = \text{accept}|R_{ij\ell t-1}, \tilde{p}_{\ell t})}{\begin{pmatrix} [1-\sigma_0(\alpha R_{ij\ell t-1}, \tilde{p}_t)] \Pr(d_{ij\ell t} = \text{accept}|R_{ij\ell t-1}, \tilde{p}_{\ell t}) \\ +[1-\sigma_0(\alpha R_{ij\ell t-1} + (1-\alpha), \tilde{p}_t)] \Pr(d_{ij\ell t} = \text{reject}|R_{ij\ell t-1}, \tilde{p}_{\ell t}) \end{pmatrix}}
\end{aligned}
$$

if $t < T_{ij\ell}$ and

$$
\begin{aligned}
&\Pr(d_{ij\ell t} = \text{accept}|R_{ij\ell t-1}, \tilde{p}_{\ell t}, \text{non-surviving}) \\
=&\frac{\sigma_0(\alpha R_{ij\ell t-1}, \tilde{p}_{\ell t}) \Pr(d_{ij\ell t} = \text{accept}|R_{ij\ell t-1}, \tilde{p}_{\ell t})}{\begin{pmatrix} \sigma_0(\alpha R_{ij\ell t-1}, \tilde{p}_t) \Pr(d_{ij\ell t} = \text{accept}|R_{ij\ell t-1}, \tilde{p}_{\ell t}) \\ +\sigma_0(\alpha R_{ij\ell t-1} + (1-\alpha), \tilde{p}_t) \Pr(d_{ij\ell t} = \text{reject}|R_{ij\ell t-1}, \tilde{p}_{\ell t}) \end{pmatrix}}
\end{aligned}
$$

if $t = T_{ij\ell}$ and the relationship is ended because of a demotion.

To speed up estimation, we loop through the common cost variance $\sigma^c \in \{0.1, 0.2, ..., 2.0\}$ in an outer loop and estimate mean transformed cost $\tilde{\mu}^c_{ia\ell}$ and transformed rent $(\eta_{ij\ell}+p_{ij\ell}+\kappa_\ell)$ for each relationships with $T_{ij\ell} \geq 30$ in an inner loop.

Step 2.2. Estimate cost parameters and extrapolate cost estimates to all lanes

We exploit our estimates $\{\widehat{\tilde{\mu}}^c_{ia\ell}\}_{T_{ij\ell}\geq 30}$ of mean transformed costs in two ways. First, we use them as the left-hand-side variable to estimate the cost parameters in Equation (9). Second, we extrapolate these cost estimates to all relationships by running a flexible polynomial regression of these estimates on observed relationship and lane characteristics,

$$\widehat{\tilde{\mu}}^c_{ia\ell} = h(\tilde{\mathbf{X}}_{ia\ell}) + \tilde{\epsilon}^c_{ia\ell}. \tag{15}$$

Together, the extrapolated transformed costs from Equation (15) and our estimates of cost parameters in Equation (9) provide estimates $(\hat{\kappa}_\ell)_\ell$ of search costs and and $(\hat{\mu}^c_{ia\ell})_{ia\ell}$ of mean operational costs in all relationships. The common cost variance $\hat{\sigma}^c$ is obtained from Step 2.1.

Recall that $\tilde{\mathbf{X}}_{ia\ell}$ includes both observable relationship-level and lane-level characteristics. The relationship-level characteristics include the log of average monthly volume and the coefficient of variation of the shipper's weekly volume. The lane-level characteristics include the average spot rate Rate_ℓ, the average spot volume $\text{Volume}^{\text{spot}}_\ell$, the average distance Distance_ℓ on lane ℓ, an indicator Tight_a of whether the relationship was formed during a tight market, and the measure Imbalance_ℓ of backhaul volume relative to forehaul volume. Note that Rate_ℓ, $\text{Volume}^{\text{spot}}_\ell$, Tight_a, and Imbalance_ℓ are equilibrium objects, which capture

the underlying differences in demand and supply factors across lanes. All remaining steps in our estimation condition on $\tilde{\mathbf{X}}_{ia\ell}$.

Step 2.3. EM-algorithm pooling all relationships

Given the cost estimates $((\hat{\kappa}_\ell)_\ell, (\hat{\mu}^c_{ia\ell})_{ia\ell}, \hat{\sigma}^c)$ from the previous steps, we flexibly estimate the conditional distribution of winning carriers' rents by a conditional mixture model. For each group $g = 1, ..., 7$ with mixing probability π_g, the component distribution of the winning carrier's rent r is a Normal distribution, whose conditional mean varies linearly with observed relationship and lane characteristics as follows

$$r|_g \sim \text{Normal}(\beta^r_g \tilde{\mathbf{X}}_{ia\ell}, \sigma^r_g).$$

For estimation, we adapt an EM algorithm by Train (2008): The M-step integrates the likelihood contribution in Equation (13) over the mixture of carrier rents. The E-step updates the mixing probabilities $\{\pi_g\}_g$, the mean-shifting parameters $\{\beta^r_g\}_g$, and variances $\{\sigma_g\}_g$ of the component distributions. To speed up the integration step, we perform linear interpolations over a grid $\{0.0, 0.1, ..., 5.0\}$ of carrier rent. That is, we calculate the likelihood contribution of each relationship at each grid point of carrier rent only once, before the EM-algorithm. The iterative procedure of the EM-algorithm only updates the weights being put on these grid points.

The output of the EM-algorithm is the distribution of the winning carrier's rent r conditional on $\tilde{\mathbf{X}}_{ia\ell}$ and on $r \geq \underline{r}_{ia\ell}$, where $\underline{r}_{ia\ell}$ is the lowest carrier rent in auctions with characteristics $\tilde{\mathbf{X}}_{ia\ell}$. We take the fifth percentile of this distribution as our estimate of $\underline{r}_{ia\ell}$.

Step 3. Estimate conditional bidding functions

With a stochastic number of bidders, $n_a \sim \text{Binomial}(N, \tilde{q})$, the distribution of the winning carrier's rent is

$$\frac{\left(1 - \tilde{q} + \tilde{q}\left[\frac{G^{\eta+p}(r|\tilde{\mathbf{X}}_{ia\ell}) - G^{\eta+p}(\underline{r}_{ia\ell}|\tilde{\mathbf{X}}_{ia\ell})}{1 - G^{\eta+p}(\underline{r}_{ia\ell}|\tilde{\mathbf{X}}_{ia\ell})}\right]\right)^N - (1-\tilde{q})^N}{1 - (1-\tilde{q})^N}.$$

This distribution, together with our estimates of $(N, \tilde{q})$ from Step 1, pins down the winning probability of a carrier with rent $r \geq \underline{r}_{ia\ell}$,

$$\left(1 - \tilde{q} + \tilde{q}\left[\frac{G^{\eta+p}(r|\tilde{\mathbf{X}}_{ia\ell}) - G^{\eta+p}(\underline{r}_{ia\ell}|\tilde{\mathbf{X}}_{ia\ell})}{1 - G^{\eta+p}(\underline{r}_{ia\ell}|\tilde{\mathbf{X}}_{ia\ell})}\right]\right)^{N-1} \equiv [\tilde{G}^{\eta+p}(r|\tilde{\mathbf{X}}_{ia\ell})]^{N-1}.$$

Our estimation of the conditional bidding functions $\mathbf{b}(\cdot|\tilde{\mathbf{X}}_{ia\ell})$ follows our identification argument closely, though with two exceptions. First, since we allow for a stochastic number

of bidders, the winning probability of a carrier with rent r is $[\tilde{G}^{\eta+p}(r|\tilde{\mathbf{X}}_{ia\ell})]^{N-1}$. Second, we select the lowest shipper rent as the level that sets the lowest match quality to zero rather than one that makes the shipper's individual rationality constraint bind at the lowest-quality relationship. We take this approach for two reasons. First, using this approach, our estimates provide lower bounds on shippers' match-specific gains; higher estimates of these gains would only strengthen our conclusion that relationships' benefits outweigh there negative externalities. Second, since the shipper' expected payoff is relatively flat in her rent for low levels of carrier rent, inverting the shipper's expected payoff at our estimate $\underline{r}_{ia\ell}$ of the lowest carrier rent would give an unreliable estimate of the lowest shipper rent.

Step 4. Estimate the joint distribution of match-specific gains

For each tuple $\tilde{\mathbf{X}}_{ia\ell}$ of observable relationship and lane characteristics, we estimate the conditional distributions $G^{\eta+p}(\cdot|\tilde{\mathbf{X}}_{ia\ell}\cup\{\, p_{ij\ell}\})$ of carrier rents and $G^{p}(\cdot|\tilde{\mathbf{X}}_{ia\ell})$ of contract rates as mixtures of conditional Normal distributions, using EM-algorithms as in Step 2.3. Given these conditional distributions and the bidding function from Step 3, we simulate for each tuple $\tilde{\mathbf{X}}_{ia\ell}$ the joint distribution of rents and contract rates $(\eta+p, p, \psi-p)$, which effectively pins down the joint distribution of match-specific gains (ψ, η).

E.2 Estimation of counterfactual institutions

For our market-level welfare analysis, we aggregate KMA-to-KMA lanes into seven lane clusters using K-means clustering on lane characteristics $\{\text{Rate}_\ell, \text{Volume}_\ell^{\text{spot}}, \text{Distance}_\ell\}$. The idea is that these characteristics would capture key differences in the underlying aggregate demand and capacity across lanes. We further classify relationships in each lane cluster into those that occured in a soft market (the year 2016) versus those that occur in a tight market (the year 2018). Thus, we have a total of 14 clusters of lanes and market phases. We perform welfare analysis separately for each of these clusters.

Estimate aggregate demand and capacity

For each cluster, the aggregate demand and capacity are estimated over a 52-week period. In what follows, we show how to recover a tuple $(L, D_t, C)_{t=1}^{52}$ of long-term demand, direct spot demand, and spot capacity for one cluster using our merged relationship and spot data.

For each week t, our sample of long-term relationships gives the total number of loads requested by the shippers, L_t^{demand}, those loads that are accepted by the primary carriers, L_t^{primary}, those that are accepted by backup carriers L_t^{backup}, and those that are rejected by all carriers and fulfilled in the spot market. Our spot data gives us the total number of spot loads $\tilde{S}_t$, which we scale to reflect the 80% aggregate share of relationship transactions (either by primary or backup carriers) across all lanes. For welfare analysis, we treat relationship

loads fulfilled by backup carriers as spot loads.

Using methods of moments, we estimate (L, C) to match relationship demand, relationship volume, and spot volume to their empirical counterparts $(L_t^{\text{demand}}, L_t^{\text{primary}}, \tilde{S}_t + L_t^{\text{backup}})$. We then calibrate spot demand $(D_t)_{t=1}^{52}$ to satisfy the market equilibrium condition in each period t,

$$\underbrace{L}_{L_t^{\text{demand}}} + D_t = \underbrace{L \int_{\tilde{p}_t}^{\infty} \hat{F}(\bar{p} - \hat{\kappa}) d\hat{\mu}(\bar{p}|\tilde{p}_t)}_{\text{LTR: } L_t^{\text{primary}}} + \underbrace{[C + L\hat{\mu}(\tilde{p}_t|\tilde{p}_t)]\hat{F}(\tilde{p}_t - \hat{\kappa})}_{\text{Spot: } \tilde{S}_t + L_t^{\text{backup}}}.$$

Here, spot rate $\tilde{p}_t$ and estimated cost parameters $(\hat{F}, \hat{\kappa})$ are the cluster-specific medians; the distribution $\mu(\cdot|\tilde{p}_t)$ of the full compensation is estimated from these cost parameters and the estimated distribution of carrier rents conditional on the median observable relationship and lane characteristics $\tilde{\mathbf{X}}_{ia\ell}$.

Our welfare calculation keeps fixed the aggregate demand and supply $(L, D_t, C)_{t=1}^{52}$and normalizes shippers' gains from having their loads shipped to zero.

Counterfactual 1. Spot-only

In a centralized spot market, long-term demand is combined with direct spot demand and all carriers in the market make up spot capacity. Thus, the equilibrium spot rate $\tilde{p}_t^1$ in each period and search cost κ^1 are pinned down by

$$L + D_t = \underbrace{(L + C)F(\tilde{p}_t^1 - \kappa^1)}_{\text{Spot: } S_t^1}, \text{ where } \kappa^1 = \kappa \left(\sum_t S_t^1 / T \right).$$

The period-t welfare of a centralized spot market is

$$W_t^1 = (L + D_t)(-\kappa^1 - \mathbf{E}[c_t | c_t \leq \tilde{p}_t^1 - \kappa^1]).$$

Counterfactual 2. Index-priced

Denote by κ^2 the equilibrium search cost under the relationship-optimal index-priced contracts. Under these contracts, any relationship with $\theta_{ij} \geq -\kappa_2$ generates non-negative surplus over spot transactions; the carrier rejects if and only if $c_t \geq \theta_{ij} + \tilde{p}_t^2$ and never serves the spot market (on the same lane). Thus, the equilibrium spot rate $\tilde{p}_t^2$ in each period and

search cost κ^2 are pinned down by

$$L + D_t = L \int_{\theta \geq -\kappa^2} F(\theta + \tilde{p}_t) d[G^\theta]^{1:N}(\theta) + \underbrace{CF(\tilde{p}_t - \kappa^2)}_{\text{Spot: } S_t^2}, \text{ where } \kappa^2 = \kappa\left(\sum_t S_t^2/T\right).$$

The period-t welfare under index-priced contracts is

$$W_t^2 = L \int_{\theta \geq -\kappa^2} F(\theta + \tilde{p}_t)(\theta - \mathbf{E}[c_t | c_t \leq \theta + \tilde{p}_t]) d[G^\theta]^{1:N}(\theta)$$
$$+ CF(\tilde{p}_t - \kappa^2)(-\kappa^2 - \mathbf{E}[c_t | c_t \leq \tilde{p}_t - \kappa^2]).$$

This aggregate welfare includes the welfare of shippers and carriers that only operate in the spot market as well as those that form relationships. For the latter group, notice that (i) the full surplus, denoted by $\text{Surplus}(\theta_{ij}|\tilde{p}_0)$, is realized, and that (ii) the shipper's surplus is precisely the fixed fee b_{ij}^0 that the winning carrier bids. For a relationship with match quality $\theta_{ij} \geq -\kappa^2$, this fee equals

$$\text{Surplus}(\theta_{ij}|\tilde{p}_0) - \frac{\int_{-\kappa^2}^{\theta_{ij}} \text{Surplus}(\theta|\tilde{p}_0) d[G^\theta]^{1:N}(\theta)}{[G^\theta]^{1:N}(\theta_{ij})}.$$

Counterfactual 3. Market-level (constrained) first-best

We demonstrate the implementation of the market-level constrained first-best welfare for one period t via an appropriate tax on index-priced contracts. Abusing notation, let L denote the continuum of primary carriers, C the continuum of spot carriers, and $\kappa(S)$ the search cost induced by spot volume S. Fix spot demand D_t in period t.

An allocation $A : j \mapsto \{\text{accept}, \text{spot}, \text{idle}\}$ maps each carrier $j \in L \cup C$ to a decision such that $A(j) \in \{\text{spot}, \text{idle}\}$ for all $j \in C$ and $\sum_{j \in L \cup C}(\mathbf{1}\{A(j) = \text{accept}\} + \mathbf{1}\{A(j) = \text{spot}\}) = L + D_t$. To realize spot volume $S \in [D_t, L + D_t]$, the constrained first-best allocation solves

$$\begin{aligned} W_t^3(S) = \max_A & \int_{j \in L} \mathbf{1}\{A(j) = \text{accept}\}(\theta_{ij} - c_{jt}) + \int_{j \in L \cup C} \mathbf{1}\{A(j) = \text{spot}\}(-c_{jt} - \kappa(S)) \\ \text{s.t.} \quad & \int_{j \in L} \mathbf{1}\{A(j) = \text{accept}\} = L + D_t - S \\ & \int_{j \in L \cup C} \mathbf{1}\{A(j) = \text{spot}\} = S. \end{aligned}$$

Let λ_1 and λ_2 be the Lagrange multipliers of the constraints on relationship and spot volumes, respectively. We can rewrite the above constrained optimization as

$$\max_A \int_{j\in L} [\mathbf{1}\{A(j)=\text{accept}\}(\theta_{ij}-\lambda_1-c_{jt})+\mathbf{1}\{A(j)=\text{spot}\}(\lambda_2-c_{jt}-\kappa(S))] + \int_{j\in C} \mathbf{1}\{A(j)=\text{spot}\}(\lambda_2-c_{jt}-\kappa(S)).$$

The solution to this problem involves each $j \in L$ choosing the maximal in $\{\theta_{ij}-\lambda_1-c_{jt}, \lambda_2-c_{jt}-\kappa(S), 0\}$ and each $j \in C$ choosing the maximal in $\{\lambda_2-c_{jt}-\kappa(S), 0\}$. That is, the constrained first-best allocation is achieved by combining index-priced contracts with a corrective tax on relationship transactions equal to λ_1; in equilibrium, spot rate is λ_2. Moreover, solving for the market-level first-best welfare reduces to choosing the optimal spot volume S.

F Other results

F.1 Sources of gains from long-term relationships

This section presents suggestive evidence on the sources of gains from long-term relationships by comparing three performance measures—acceptance rate, on-time delivery, and wait time—across carriers' asset ownership and relationship status. Acceptance rate captures carriers' *cooperation*. On-time delivery captures carriers' *quality of service*. Wait time, which is the time it takes carriers to wait, navigate, and load shipments at shippers' facilities, captures shipper-carrier *coordination*.

Table 6 presents the estimation results of the following regression

$$\text{Performance}_{ij\ell t} = \sum_c \alpha_c \mathbf{1}\{\text{carrier } j\text{'s type is } c\} + \alpha_{\text{spot}}(\tilde{p}_{\ell t} - p_{ij\ell}) + \text{controls} + \epsilon_{ij\ell t}.$$

Here, $\text{Performance}_{ij\ell t}$ is one of our three performance measures, and carrier's type captures both asset ownership (asset-owner or broker) and relationship status (primary or backup). The (OLS) specifications capture the differences in the performance between primary and backup carriers that could arise from two channels: the selection of better relationships via auctions and the gains from repeated interactions within these relationships. The (FE) specifications include shipper-carrier-lane fixed-effects to isolate the second channel.

Comparing the performance of asset-owner carriers and that of brokers in the (OLS) specification, we find suggestive evidence of match-specificity. Specifically, primary asset-owner carriers perform substantially better than primary brokers in on-time delivery (by

Table 6: Comparison of different groups of carriers across three performance measures

	Accept (0/1)		On-time (0/1)		Wait time (hours)	
	(OLS)	(FE)	(OLS)	(FE)	(OLS)	(FE)
Asset owners						
Primary	0.732	0.143	0.719	0.0114	1.514	-0.00617
	(0.00132)	(0.00109)	(0.00201)	(0.00244)	(0.0120)	(0.0152)
Backup	0.352		0.657		1.703	
	(0.00138)		(0.00237)		(0.0140)	
Brokers						
Primary	0.833	0.235	0.611	-0.00402	2.035	-0.0637
	(0.00144)	(0.00146)	(0.00219)	(0.00293)	(0.0130)	(0.0184)
Backup	0.375		0.613		2.105	
	(0.00147)		(0.00261)		(0.0155)	
spot premium	-0.185	-0.178	0.0288	-0.0230	-0.253	0.0149
	(0.000603)	(0.000825)	(0.00118)	(0.00185)	(0.00701)	(0.0116)
N	2320132	2313178	874916	850760	847218	840656

Note: (OLS): controls include the frequency and inconsistency of load timing and year dummies. (FE): controls include year dummies. The coefficient estimates for primary capture the difference between being primary versus backup for the respective type of carrier (asset owner or broker). On-time delivery and wait time are only measured for accepted shipments.

10.8 pp) and wait time (by 31 minutes). Since a broker has access to a large pool of carriers, a shipper requesting shipments via a broker ends up interacting with multiple asset-owner carriers. Thus, the lower quality and coordination of brokers speak to the potential match-specificity in the relationships between shippers and asset-owner carriers.

Comparing the performance of primary carriers and that of backup carriers in the (OLS) specification, we find suggestive evidence of the auctions' role in selecting good matches. Specifically, we find that primary asset-owner carriers perform substantially better than their backup counterparts, with the former group achieving 6.2 pp improvement in on-time delivery and 11-minute reduction in wait time. The same pattern does not hold when comparing primary and backup brokers.

Finally, comparing the (FE) results of asset-owner carriers across the three performance measures, we find that only acceptance rate meaningfully (i) changes with carriers' primary status and (ii) responds to spot premiums. These findings support our decision to focus on carriers' accept/reject decisions as their key decision margin and on spot temptation as the source of carriers' *moral hazard*. Other performance gains, such as on-time delivery and low wait time, could be interpreted as relationships' *intrinsic gains* in our framework.

Table 7: Estimates of shippers' punishment scheme

	(GMM)		(IV Probit)	
Demotion probability	Estimate	95% CI	Estimate	95% CI
R_{t-1}	0.0158	$(0.0058, 0.0662)$	0.616	$(0.319, 1.853)$
R_{t-1} × Frequency	0.0227	$(0.0128, 0.0687)$	0.584	$(0.320, 1.610)$
R_{t-1} × Inconsistency	−0.0220	$(-0.0939, 0.0794)$	0.468	$(-0.240, 2.045)$
Frequency	−0.0207	$(-0.0306, -0.0131)$	−0.650	$(-0.864, -0.464)$
Inconsistency	0.0764	$(0.0338, 0.1089)$	0.171	$(-0.663, 0.528)$
Spot premium	−0.00551	$(-0.0081, -0.0037)$	−0.156	$(-0.232, -0.111)$

Note: The common decay parameter α and initial rejection index R_0 are estimated using GMM; $\hat{\alpha} = 0.985$ (95% CI =$(0.956, 0.995)$); $\hat{R}_0 = 0.8$ (95% CI = $(0.0, 1.0)$). Frequency is the log of average monthly volume; Inconsistency is the average coefficient of variation of weekly volume within a month. The confidence intervals are constructed from 50 full-estimation bootstraps, with sampling at the auction level. $N = 572274$ observations for both specifications.

F.2 Shipper's punishment scheme and other instrumental objects

Table 7 presents our estimates of shippers' punishment scheme. The coefficient of the rejection index R_{t-1} is positive and significant. To interpret the magnitude of this coefficient, we simulate two sets of relationships, one with an initial rejection and one with an initial acceptance. We find that a carrier's initial rejection instead of an acceptance would reduce the expected number of requests he would receive by 3%. This suggests that the shipper's punishment scheme is soft but generates meaningful incentives.

Moreover, we find that when the current spot rate ($\tilde{p}_t$) is high or if the shipper has large volume on a lane, that is, when the relationship is more valuable to the shipper, demotion probability is lower. Finally, the daily decay parameter α on the carrier's past rejections is close to one, suggesting that a rejection of the carrier affects the continuation probability of his relationship in many periods. These findings on shippers' punishment scheme are consistent with findings in our companion paper Harris and Nguyen (2024).

F.3 Distributional consequences

Figure 12 presents estimated changes in the average per-period payoff ($/mile) of shippers and carriers who manage to form relationships and those only operating on the spot market under the *spot-only* and *index-priced* counterfactuals, relative to the current institution with fixed-rate contracts in long-term relationships.

Figure 12: Distributional consequences by different groups of shippers and carriers

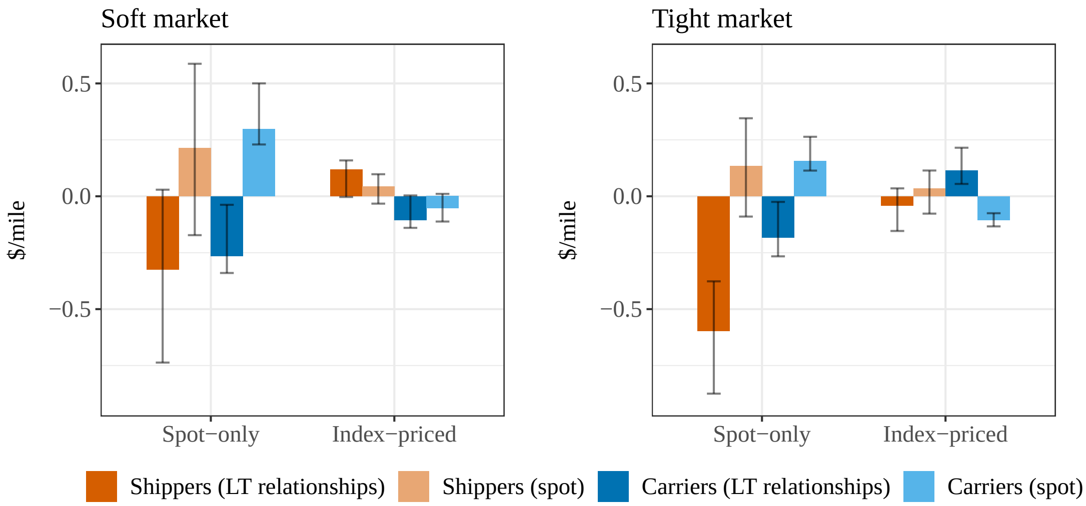

Note: Shippers' benefits from having their loads shipped are normalized to zero.

Spot-only. Our findings suggest that eliminating relationships would substantially hurt shippers and carriers who would manage to form relationships in the current institution. However, spot carriers would unambiguously benefit from the *spot-only* counterfactual, due to both (i) an increase in demand for spot loads and (ii) a reduction in search costs for spot loads.

Index-priced. While this counterfactual would barely change aggregate welfare from the current institution, we find substantial distributional consequences when breaking down welfare changes by agents' types. First, we find that the *index-priced* counterfactual would favor shippers over carriers, with the largest welfare loss borne by carriers operating only on the spot market. These carriers tend to be small owner operators, who might exit the market if faced with such loss.

A more subtle finding is that the global adoption of the relationship-optimal *index-priced* contracts in place of the current fixed-rate contracts would not guarantee welfare gains for relationship participants. Moreover, market conditions would affect how such gains (if any) are split. Specifically, we find that relationship carriers would incur substantial welfare loss from the *index-priced* counterfactual in a soft market, when their bargaining position is low. The reason is that a thinner and less efficient spot market in the *index-priced* counterfactual would worsen carriers' outside option, directly hurting them when exercising this option and indirectly hurting their bargaining outcomes at the auction stage.