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

Adam Harris and Thi Mai Anh Nguyen*

May 13, 2025

Abstract

The prevalence of relationships suggest that they offer significant benefits to participants over spot transactions. However, some of these benefits might be self-enforcing as spot markets feature scale efficiencies. We quantify the optimal balance between relationship and spot transactions in the US truckload freight industry, where long-term relationships with fixed-rate contracts dominate. We find that the incompleteness of fixed-rate contracts limits relationship-level surplus, acting as a corrective tax on relationship transactions. Overall, the current institution achieves 92% of the market-level first-best surplus, outperforming both the relationship-optimal and spot-optimal scenarios. A marginal increase in spot transactions would improve market-level welfare.

1 Introduction

Long-term relationships and spot transactions are ubiquitous—and typically competing—features of the economy (Baker, Gibbons and Murphy, 2002). While spot transactions present an outside option to relationships, making relational incentives harder to enforce, the formation and high performance of relationships could result in a thinner and less efficient spot market (Kranton, 1996). These counteracting forces mean that the observed balance between relationship and spot transactions may deviate from the social optimum.

*Harris: Cornell University, 109 Tower Rd, Office 436, Ithaca, NY 14853, adamharris@cornell.edu. Nguyen: New York University, 19 West 4th Street, New York, NY 10012, tmn285@mit.edu. 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.

Determining the direction of the deviation from this optimum—and therefore whether policymakers should encourage more relationships or more spot transactions—is an important empirical question. In recent years, digital spot platforms have disrupted various markets (e.g., ride-hailing, delivery) but have not made significant inroads in markets in which relationships have traditionally played a central role.[1] In such markets, do relationships represent a barrier to a potentially more efficient market structure, or should they evolve to better insulate themselves from the competitive pressure of spot transactions?

We study these questions in an economically important industry—US for-hire truckload freight—where the interplay between long-term relationships and spot transactions is a central feature. In this setting, 80% of transactions occur under long-term contracts that fix rates but not volume, as compared to 20% by spot arrangements.[2] While the incompleteness of fixed-rate contracts exposes relationships to spot temptation, resolving this contractual issue within relationships may exacerbate the inefficiency of the already-thin spot market. We develop an empirical framework to quantify both of these economic forces and evaluate the current market structure against alternative splits between relationship and spot transactions. We find that social welfare would benefit from a marginal—but not a global—shift toward spot transactions. This reflects both the substantial potential for scale efficiency in the spot market and the substantial intrinsic benefits of relationships. Moreover, these findings suggest that digital spot platforms that enhance spot market efficiency and complement, rather than replace, relationships should be encouraged.[3]

In US for-hire truckload freight, a firm with demand for transportation service ("shipper") on an origin-destination pair ("lane") hires a trucking company ("carrier") for that service. Two empirical advantages of this setting allow us to quantify both the rich dynamics within shipper-carrier relationships and the intricate interplay between these relationships and the spot market. 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 across all of their contracted lanes, we observe the details of all of these shippers' relationships, including their offers, carriers' responses, and the status of each relationship at each point in time.[4] Second, variation in transportation demand results

[1]Examples of such relationship-based markets can be found in supply chains (e.g., Boudreau, Cajal-Grossi and Macchiavello, 2023), finance (e.g., Allen and Wittwer, 2023), and most labor contexts (BLS, 2024).

[2]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, Darmouni and Essig Aberg (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. Our analysis takes this fixed-price feature as a given contractual friction, focusing instead on its welfare implications.

[3]Providers of digital trucking solutions seem to share this view. For example, Uber Freight offers digital spot option that can be integrated into shippers' TMS for contract freight (Berdinis, 2024).

[4]Since each of these shippers interacts with multiple carriers across a broad set of lanes, our data exhibits

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 (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 two data patterns that highlight the key tradeoff between the intrinsic benefits of relationships and scale efficiency of the spot market. First, we find substantial within-auction variation in contract rates across carriers even after controlling for potential differences in carriers' costs. Moreover, we find that shippers do not necessarily select the carrier with the lowest contract rate as the auction winner ("primary carrier"), doing so in only two thirds of the auctions. These observations suggest that both carriers and shippers care about non-price factors in long-term relationships. Second, 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 institutional details of our setting and the data patterns described above. At the relationship level, each relationship comprises two stages. In the first stage, the shipper's and carriers' match-specific gains are drawn from a distribution; 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. The strategy space of this repeated game is informed by our findings in Harris and Nguyen (2025) on the nature of shipper-carrier interactions within relationships. In each period of this game, the shipper decides based on the carrier's past rejections 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. At the market level, the spot market absorbs both direct spot demand and unfulfilled relationship demand, pinning down the equilibrium spot rate and volume. Moreover, search costs for spot loads might scale (inversely) with spot market volume. Thus, relationships exert both pecuniary and non-pecuniary externalities on the spot market.

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 carriers' optimal response to spot conditions in the repeated game. To build intuition, consider the case in which carriers only face static incentives. When spot rates are sufficiently low, a carrier accepts as long as his costs are low; when spot rates are sufficiently high, he never

national geographical coverage.

accepts. At the critical spot rate at which acceptance drops to zero, the carrier is indifferent between *accept* and *spot*. Thus, this critical spot rate pins down the carrier's non-price gain from the relationship, and the carrier's acceptance to the left of this critical point reveals information about his search and operational costs. Our formal identification argument generalizes this intuition to allow for carriers' dynamic incentives, which arise as shippers punish carriers' rejections with an increased likelihood of relationship termination. Finally, to identify the causal link between search costs and spot market volume, we use as a demand shifter the predicted trade flows between US states from Caliendo et al. (2018).

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, Perrigne and Vuong (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. These premiums are substantial; for context, consider that the average spot rate over our sample period is \$1.87/mile. 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 spot market thickness substantially reduces 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 performance and share of relationships. At the relationship level, we find that the current relationships generate substantial surplus of \$0.64/mile to their participants over spot transactions. However, this realized surplus represents only 64% of the relationship-level first-best surplus. The latter could be achieved through *index-priced* contracts that are pegged one-to-one to spot rates. These relationship-level findings explain the prevalence of relationships in the truckload setting while suggesting strong incentives for relationships' participants to switch to more flexible contracts.

At the market level, we evaluate the optimal balance between relationship and spot transactions while keeping fixed the aggregate demand and capacity. In our framework, a higher share of relationship transactions realizes more relationship benefits but also leads to higher search costs. Directly, higher search costs burden spot carriers, and indirectly, create stronger incentives to avoid search costs, rationalizing more relationship transactions with negative intrinsic gains or high operational costs. To highlight the welfare implications of the split between relationship and spot transactions, we trace out the highest achievable welfare at each level of the spot market's share of total market volume. Each point of this constrained first-best welfare curve can be implemented by pairing index-priced contracts with an appropriate tax on relationship transactions.

Our counterfactual analysis yields three main findings. First, relationships remain dominant in the market-level first-best allocation, and the *spot-only* scenario is the lowest point on the constrained first-best curve. This finding underscores the critical role of relationships in our setting. Second, the current institution with fixed-rate contracts realizes 92% of the market-level first-best surplus, marginally outperforming the zero-tax *index-priced* scenario. This suggests that contractual frictions within current relationships hurt social welfare less than the lack of scale in the current spot market. Third, and relatedly, marginally increasing the share of spot transactions would improve social welfare. Our takeaway from this finding is not that the government should tax relationship transactions—which is impractical. Rather, this finding highlights the benefits of encouraging technological advances (e.g., algorithmic matching and pricing) that enhance spot market efficiency; in addition to their direct benefits to spot participants, such improvements would shift the equilibrium split between relationship and spot transactions toward the market-level first-best.

Related literature. Our paper contributes to three main strands of literature. *First*, we contribute to the literature on long-term relationships.[5] Conceptually, the theoretical forces at the heart of our analysis are closely related to those described by Kranton (1996). Kranton posited that the two-way crowding-out effects between relationships and spot markets can give rise to multiple equilibria—some in which relationships predominate and others in which spot predominates—with no guarantee that the socially optimal equilibrium will be selected. This theoretical result begs an empirical question: In markets where relationships and spot markets coexist, can we determine whether there is too much or too little spot activity relative to the social optimum? Our main contribution to this literature is an empirical framework that allows us to quantify (i) the market-level tradeoffs between relationships

[5]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, Macchiavello and Noguera, 2023), reputation building (Macchiavello and Morjaria, 2015), or relational adaptation (Barron et al., 2020). Relationships may terminate (Macchiavello and Morjaria, 2015) or restructure (Gil, Kim and Zanarone, 2021) in the face of large shocks, and can be hampered by competition (Macchiavello and Morjaria, 2021).

and spot and (ii) the welfare effects of shifting the market-level split between these two modes of transaction.

Our approach to quantifying relationship primitives builds upon previous insights. 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 lead 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. We then recover shipper primitives from the relationship formation process by formulating it as an auction on rents. This auction-based approach (Guerre, Perrigne and Vuong, 2000; Asker and Cantillon, 2008) 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 ideas that (i) the stickiness of long-term bilateral interactions can result in market inefficiencies and that (ii) such stickiness can arise endogenously are not unique to our setting. In the liquefied natural gas market, for example, Zahur (2022) shows that contract lock-in can lead to ex-post allocative inefficiencies. Similarly, in the context of containership leasing contracts, Garcia-Osipenko, Vreugdenhil and Zahur (2025) show that private parties might settle on contracts with socially inefficient duration.[6] While we study conceptually similar spillovers of long-term contracts, the underlying strategic behaviors differ due to the nature of contracting in our setting: the LNG and containership settings feature full contractual commitments, whereas ours features incomplete contracts sustained by relational incentives. This means that market dynamics in our setting are shaped by strategic behaviors not only at the formation stage, as in these other papers, but also throughout the course of relationships.

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, Lizzeri and Salz, 2019; Brancaccio, Kalouptsidi and Papageorgiou, 2020; Buchholz, 2022), a problem

[6]These contract choices, in turn, affects market thickness/thinness—and therefore efficiency. This idea is related to liquidity externalities studied in financial markets, where a common solution is to centralize transactions on 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 because dealer competition on the platform is low. 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).

that might be mitigated by large spot platforms with economies of density and pricing instruments.[7] Our paper is similar to these other papers in examining scale efficiency of the spot market but differs in our emphasis on the role of long-term relationships. A novel insight of our paper is that while existing relationships are valuable to participating parties, they might also serve as a market-level barrier to spot platform adoption.[8]

Finally, 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).[9] 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. This paper also builds upon our companion paper Harris and Nguyen (2025), which used the same data to provide descriptive evidence on the interactions within *individual shipper-carrier relationships*. While those findings inform how we model individual relationships in the current paper, we go beyond this previous work in many ways. Most notable is the focus of our current work on the *market-level* interactions between relationships and spot transactions. We raise a general question—should relationships play a more or less prominent role?—and answer it empirically in the context of US truckload freight. Specifically, we develop a new structural framework to quantify the intrinsic value of relationships and the scale efficiency of the spot market, assessing their tradeoff under different market institutions.[10]

[7] These large platforms may also extract most of the surplus they create (e.g., Rosaia, 2023; Brancaccio et al., 2023). Furthermore, platform efficiency depends on the specifics of its design, such as the sophistication of the pricing rule (e.g., Castillo, 2023) or the mechanism for eliciting users' heterogeneous preferences (e.g., Gaineddenova, 2022; Buchholz et al., 2024).

[8] Digital platforms have great potential, not only because they might leverage technology to facilitate matching and pricing, but also because they can exploit various contractual and market designs. For example, Faltings (2025) studies the efficiency of flexible contracts—which allow carriers to renege on accepted offers at some reputational penalty—used by a pioneering digital brokerage. Despite this promise, digital spot platforms account for a very small share of all trucking transactions. For example, in 2024Q4, Uber Freight (the largest such platform) reported gross freight bookings of \$1.273 billion (Uber Technologies, 2025), representing less than 1% of total US for-hire trucking revenue.

[9] 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, Caplice and Sheffi, 2020), factors that affect carriers' value of relationships (Acocella, Caplice and Sheffi, 2022*b*), and bilateral improvements from index pricing (Acocella, Caplice and Sheffi, 2022*a*).

[10] Given our primary interest in this relationship-spot tradeoff, we 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 for studying 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) 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.[11] 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.[12] 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 optimally combining shipments to fill up their trucks, a concern relevant to 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).[13] To further reduce heterogeneity, our analysis focuses on dry-van (as opposed to refrigerated, flatbed, or tanker) services and on long-haul lanes of at least 250 miles.

Shippers and carriers in the US for-hire truckload industry engage in two 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*b*). The next subsections provide institutional details for each of these market components.

[11] This statistic is calculated using data from the Bureau of Transportation Statistics.

[12] 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 dense network of truck movements that allows for efficient routing.

[13] 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.

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 compensates carriers for changes in fuel costs.[14] Based on these bids, the shipper then chooses a *primary carrier*, as well as a set of backup carriers who may receive requests for any loads that the primary carrier rejects.[15]

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.[16] Our analysis focuses on the relationships between shippers and their primary carriers.

A crucial 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.

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. Our previous paper, Harris and Nguyen (2025), presents empirical evidence on the dynamics within shipper-carrier relationships. Three key insights from that paper will inform our mod-

[14] 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.

[15] 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.

[16] 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 3 in Appendix C provides an example of the sequential request process.

eling choices in this paper. First, we documented that carriers behave opportunistically—tending to reject loads within the relationship more frequently when spot rates are high relative to contract rates—while shippers do not. Second, shippers deter such carrier opportunism by employing a termination strategy. In particular, frequent rejections by the primary carrier increase the likelihood that he is permanently demoted from the primary position and loses substantial future volume from the shipper.[17] Third, this termination strategy may not condition on the entire history of the shipper-carrier relationships across the multiple lanes on which they interact. Rather, we found that when a shipper and an asset-owner carrier transact on multiple lanes, a carrier's rejection on lane ℓ may be punished by a demotion on lane ℓ, but will not be punished by a demotion on some other lane ℓ'.[18] This finding informs the unit of analysis in this paper: we consider relationships at the shipper-carrier-lane level, rather than the shipper-carrier level.

2.3 Spot transactions

Spot transactions occur through multiple channels, all of which involve some degree of 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 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 (CHR, 2015-2019). 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, destina-

[17] As we argued in the previous paper, the fact that demotions are permanent rules out the possibility that learning, rather than punishment, explains the dynamics of shipper-carrier relationships.

[18] Harris and Nguyen (2025) finds evidence that, when deciding whether to terminate brokers' primary status, shippers consider brokers' acceptances and rejections across all lanes. The current paper focuses exclusively on asset-owner carriers.

tion, distance, and activity date.[19] 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.[20]

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. In Appendix D, we discuss the representativeness of this sample. While we find that large carriers are overrepresented in our sample relative to the population of all carriers, this is consistent with the fact that large carriers are more likely to form relationships (Holm, 2020*a*). Moreover, comparing the performance of carriers in our sample across the size distribution, we find little heterogeneity, suggesting that any overrepresentation of large carriers is unlikely to bias our estimates.

Overall, 70.2% of loads are accepted by primary carriers, 19.7% by backup carriers, and 10.1% of loads 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 their relationship status, treating loads fulfilled by backup carriers as spot arrangements.[21]

We obtain spot data from DAT Solutions (DAT, 2015-2019), the dominant freight marketplace platform in the US and the leading vendor of spot market data. DAT partitions 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 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 average $1.87/mile and $1.83/mile, respectively. Both exhibit large and persistent dif-

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

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

[21] Appendix E.1 compares the performance of asset-owners versus brokers and primary versus backup carriers on three dimensions: load acceptance, on-time delivery, and detention 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 a different cost structure and a different relationship with shippers than do asset owners. Backup asset-owner carriers also perform worse than primary asset-owner carriers, with performance that, in many respects, resembles that of backup brokers.

ferences across lanes, with lane fixed-effects having a 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.[22] Table 4 in Appendix C provides key summary statistics of our merged relationship and spot data.

3.2 Key facts

We begin by establishing two empirical facts that speak to the market-level tradeoff between long-term relationships and spot transactions: (i) that relationships' match-specific gains are large and heterogeneous and (ii) that the spot market has potential scale efficiencies.

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. First, examining bidding behavior, the left panel of Figure 1 plots the distribution of contract prices (carriers' bids) after residualizing auction fixed effects and carrier-lane fixed effects.[23] The latter set of fixed effects absorbs potential differences in carriers' costs and regional specialization. We observe substantial variation in these residualized bids, with a standard deviation of \$0.22/mile. Such variation suggests that match-specific gains from relationships play a critical role for carriers. Second, turning to shippers' selection of primary carriers, the right panel of Figure 1 plots the difference between the bids of the primary carrier and the lowest-bid backup carrier across auctions. In one-third of the auctions, the primary carrier is not the lowest bidder, and in these instances, the median primary-backup price gap is \$0.16/mile. By revealed preference then, shippers care about non-price factors when forming relationships with carriers.

Relationships' match-specific gains could arise from a number of different sources. Appendix E.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

[22] See Figure 1 in Harris and Nguyen (2025) for a time series of average contract rates and spot rates.
[23] For this purpose, lanes are broadly defined: Census region of origin $\times$ Census region of destination.

Figure 1: Suggestive evidence that non-price factors matter in relationships

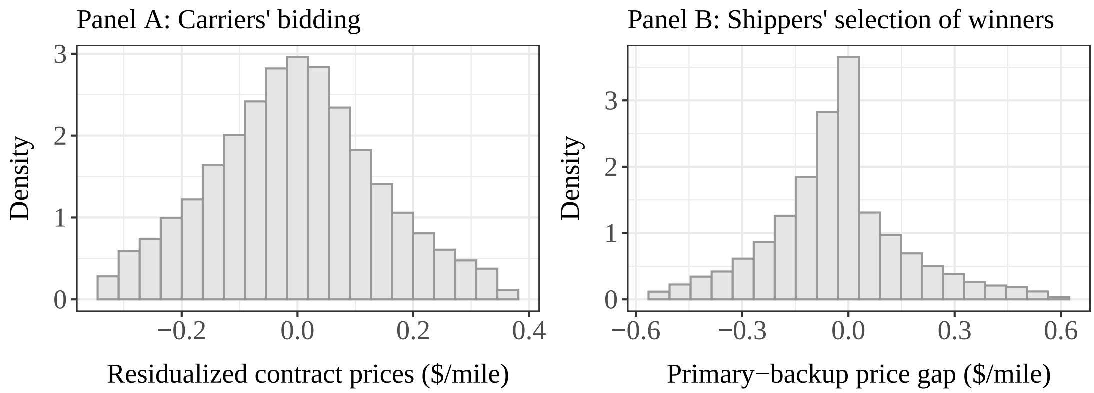

Note: Panel A: Observed contract prices are residualized on auction fixed effects and carrier-origin-destination fixed effects, with origin and destination defined at the level of census regions. The former set of fixed effects allows us to focus on within-auction variation in carriers' bids; the latter set of fixed effects further residualizes variation due to differences in carriers' costs. The histogram is at the bid level. Panel B: The histogram is at the auction level, plotting the difference between the winning (primary) bid and the lowest losing (backup) bid.

at facilities relative to backup carriers.[24] 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. Relationships can also differ in other dimensions, such as the consistency in the timing of shipments, insurance policies, billing practices, and en-route communication.

Fact 2. Spot transactions take a larger share on higher-volume lanes

Next, 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:

$$\ln(\text{Volume}^{\text{spot}}_{ss'}/\text{Volume}^{LT}_{ss'}) = \beta_0 + \beta_1 \ln(\text{Volume}^{\text{total}}_{ss'}) + \text{controls}_{ss'} + \epsilon_{ss'}. \tag{1}$$

On each state-to-state lane, $\text{Volume}^{\text{spot}}_{ss'}$ is the average number of load posts per week in the spot market, $\text{Volume}^{\text{LT}}_{ss'}$ is the average number of loads accepted per week 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

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

Table 1: 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.

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 1 presents our estimates of equation (1). 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 predicted state-to-state trade flows (Caliendo et al., 2018), and (IV2) uses origin and destination population densities.[25] 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.

[25] 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.

Figure 2: Model overview

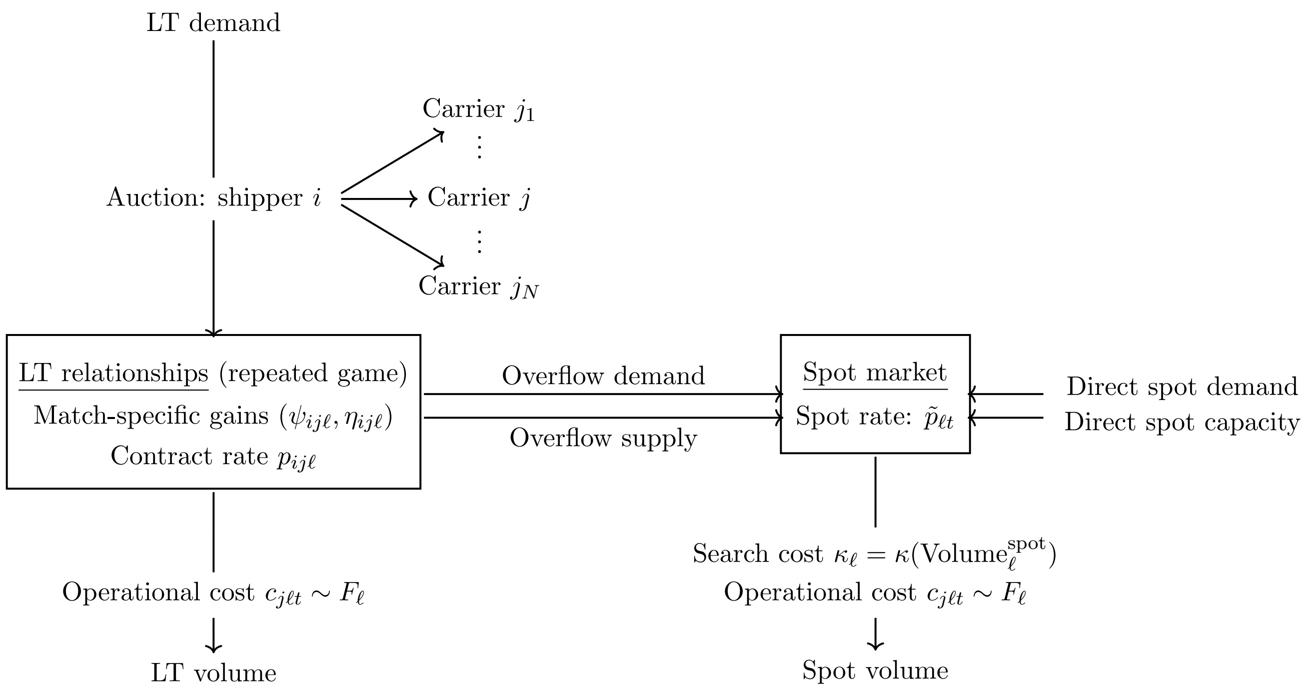

4 Model

With the goal of quantifying the market-level value of long-term relationships, we develop a model that captures both (i) the dynamics of shipper-carrier interactions within relationships and (ii) the market-level interactions between these relationships and the spot market. We will describe these two components in turn. First, we describe equilibrium behaviors within individual relationships, taking the market processes as given. Second, we describe the equilibrium conditions that give rise to these market processes. Figure 2 provides an overview of our model.

4.1 Individual relationships

4.1.1 Model primitives: timing, actions, and payoffs

A 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 the formation stage, shipper i holds an auction a to select a primary carrier. At the beginning of this auction, the shipper announces the expected frequency of interactions $\delta_{ia\ell} = \delta$. 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.[26] Then, each carrier j' proposes a contract rate $p_{ij'\ell}$. Finally, shipper i either chooses to form a relationship with some carrier $j \in J_a$ (called the primary carrier) or chooses to always go to the spot market.

If shipper i chooses j as the primary carrier, they then interact repeatedly until the shipper terminates the relationship. In each period $t \geq 1$, there is a new spot rate $\tilde{p}_{\ell t}$ and a draw $c_{j\ell t} \sim \mathcal{F}_\ell$ of carrier j's operational cost. The spot rate is observed by both the shipper and carrier; the operational cost is privately observed by the carrier. Given the observed spot rate, the shipper decides whether to maintain or end her relationship with the carrier. If the relationship is ended, both the shipper and carrier resort to the spot market for all future transactions. If, on the other hand, the relationship is maintained, the shipper offers a load to the carrier at price $p_{ij\ell}$. Observing the spot rate and his costs, the carrier chooses $d_t \in \{\text{accept}, \text{spot}, \text{idle}\}$: $d_t = \text{accept}$ 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. The shipper only observes whether the carrier accepts or rejects her load—that is, whether $d_t = \text{accept}$ or $d_t \in \{\text{spot}, \text{idle}\}$.

The instantaneous payoffs of the shipper and the carrier in period t of the repeated game depend on their actions as follows. 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; otherwise, the shipper gets service from the spot market and pays the spot rate, $u_{ij\ell t} = -\tilde{p}_{\ell t}$. 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, $v_{ij\ell t} = 0$. 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 established in the auction; and κ_ℓ is the carrier's cost of searching for a spot load on lane ℓ. 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.

4.1.2 Equilibrium behaviors

We derive the equilibrium behaviors of individual shippers and carriers by working backwards. First, we derive carriers' optimal play 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.

[26] This informational assumption matches 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, including their contracts with other shippers on other lanes.

Taking these expected payoffs as input, we then derive carriers' optimal bidding and shippers' optimal selection of primary carriers in auctions for fixed-rate contracts, which we transform into auctions over per-transaction rents. For tractability, we make the following assumption on agents' beliefs about the market processes and on shippers' strategy.

Assumption 1. *The following hold for all shippers and carriers.*

(i) (Agents' beliefs about market processes) Shippers and carriers take the market processes as given. In particular, they perceive that the spot rate $\tilde{p}_{\ell t}$ on lane ℓ follows an AR(1) process $\mathcal{P}_\ell$ and that search cost $\kappa_\ell = \kappa(Volume_\ell^{spot})$ is a function of the average spot volume on lane ℓ.

(ii) (Shippers' punishment scheme) Each shipper conditions relationship termination on an index summarizing her primary carrier's past rejections, defined recursively by $I_R : (R_{t-1}, d_t) \mapsto R_t = \alpha R_{t-1} + (1-\alpha)\mathbf{1}\{d_t \neq \text{accept}\}$, with known α and R_0. Moreover, the probability $\sigma_0 : (R_{t-1}, \tilde{p}_t) \mapsto [0,1]$ of the shipper terminating her relationship with the carrier is a function of only the carrier's rejection index and current spot rate.[27]

These assumptions have two key advantages. First, they imply that relationships evolve with a Markov state $(R_{t-1}, \tilde{p}_{\ell t})$. This is the minimal state space that both allows carriers' past rejections to influence relationship termination and captures the autocorrelated nature of the spot process. Second, the assumption that shippers' punishment schemes do not depend on match-specific gains simplifies our identification argument. An alternative approach that allowed the punishment scheme to vary with match-specific gains and exploited its optimality for identification of these gains would be sensitive to modeling assumptions on shippers' commitment power and strategy space.[28] Thus, we treat shippers' punishment scheme as given and, rather than attempting to rationalize its optimality, focus instead on the optimality of shippers' behaviors on other decision margins. Similarly, our counterfactual analyses do not require taking a stance on how shippers choose punishment schemes, as we focus on counterfactual scenarios that eliminate contractual frictions—and thus the need for such schemes.

[27] For estimation, we will allow punishment schemes to depend on observables of relationships.

[28] In settings where principals can commit to dynamic contracts that feature flexible transfers, identification follows from the interim first-order conditions (e.g., Brugues, 2020). This approach does not apply to our setting due to the lack of flexible transfers. The alternative approach that assumes no commitment power from shippers faces other challenges. First, one set of parameters can lead to multiple equilibria—a classic prediction in the literature on repeated games (Fudenberg and Maskin, 1986). Second, classic results do not readily explain the empirical observation that on-path demotions of primary carriers occur and are almost always permanent (Harris and Nguyen, 2025). In this previous paper, we also found evidence supporting the presence of shippers' commitment in relationships with large asset-based carriers.

Repeated game. In each period of the repeated game, the carrier weighs his current payoff against how his decision would affect the probability of relationship termination in future periods. Denote by $\sigma_{ij\ell} : (R_{t-1}, \tilde{p}_{\ell t}) \mapsto \{\text{accept, spot, idle}\}$ the carrier's optimal Markov strategy that breaks ties in favor of acceptance.

We can show that the carrier's optimal decision in each period t reduces to choosing the maximum among:

$$\begin{cases} \bar{p}_{ij\ell}(R_{t-1}, \tilde{p}_{\ell t}) - \tilde{c}_{j\ell t} & \text{, if } d_t = \text{accept} \\ \tilde{p}_{\ell t} - \tilde{c}_{j\ell t} & \text{, if } d_t = \text{spot} \\ 0 & \text{, if } d_t = \text{idle,} \end{cases} \tag{2}$$

where $\tilde{c}_{j\ell t} \equiv c_{j\ell t} + \kappa_\ell$ is the *transformed cost*—the carrier's cost of servicing a spot load—and

$$\begin{aligned} &\bar{p}_{ij\ell}(R_{t-1}, \tilde{p}_{\ell t}) \\ &\equiv \underbrace{\eta_{ij\ell} + p_{ij\ell} + \kappa_\ell}_{\text{relationship premium}} + \frac{\delta}{1-\delta} \underbrace{\left(V_{ij\ell}(\alpha R_{t-1}, \tilde{p}_{\ell t}) - V_{ij\ell}(\alpha R_{t-1} + (1-\alpha), \tilde{p}_{\ell t}\right)}_{\text{dynamic compensation}} \end{aligned} \tag{3}$$

is the carrier's *full compensation*—how much he is compensated for an acceptance. Here, $V_{ij\ell}(R_t, \tilde{p}_{\ell t})$ is the expected payoff of the carrier at the beginning of the next period following state $(R_t, \tilde{p}_{\ell t})$ at the end of period t. Together, equations (2) and (3) characterize optimal carrier decision-making and highlight how the benefits of accepting a relationship load versus taking a spot load shape carriers' decisions. When accepting a relationship load, the carrier's direct benefits include his per-transaction rent and his savings on search costs. Further, accepting a relationship load yields dynamic compensation equal to the difference in expected payoffs were he to accept versus were he to reject the offered load. This difference is positive if the shipper punishes a higher rejection rate with a higher likelihood of demotion.

Together, the shipper's punishment scheme σ_0 and the carrier's optimal strategy $\sigma_{ij\ell}$ determine the expected payoffs of each relationship. Following Markov state $(R_{t-1}, \tilde{p}_{\ell t-1})$, the expected payoffs of shipper i and carrier j on lane ℓ are respectively as follows:

$$\begin{aligned} &U_{ij\ell}(R_{t-1}, \tilde{p}_{\ell t-1}) \\ =&\mathbf{E}_{F_\ell, \mathcal{P}_\ell} \left[\begin{array}{l} (1 - \sigma_0(R_{t-1}, \tilde{p}_{\ell t})) \\ \times \left(\begin{array}{c} \mathbf{1}\{\sigma_{ij\ell}(R_{t-1}, \tilde{p}_{\ell t}) = \text{accept}\} \left(\delta(\psi_{ij\ell} - p_{ij\ell}) + (1-\delta) U_{ij\ell}(\alpha R_{t-1}, \tilde{p}_\ell)\right) \\ + \mathbf{1}\{\sigma_{ij\ell}(R_{t-1}, \tilde{p}_{\ell t}) \neq \text{accept}\} \left(-\delta \tilde{p}_{\ell t} + (1-\delta) U_{ij\ell}(\alpha R_{t-1} + 1 - \alpha, \tilde{p}_\ell)\right) \end{array} \right) \\ + \sigma_0(R_{t-1}, \tilde{p}_{\ell t}) \underline{U}(\tilde{p}_{\ell t}) \end{array} \middle| \tilde{p}_{\ell t-1} \right] \end{aligned}$$

and

$$
\begin{aligned}
&V_{ij\ell}(R_{t-1},\tilde{p}_{\ell t-1})\\
=&\mathbf{E}_{F_\ell,\mathcal{P}_\ell}\left[\left.\begin{array}{l}(1-\sigma_0(R_{t-1},\tilde{p}_{\ell t}))\\ \times\left(\begin{array}{c}(1-\delta)\max\{\bar{p}(R_{t-1},\tilde{p}_{\ell t})-\tilde{c}_{j\ell t},\tilde{p}_{\ell t}-\tilde{c}_{j\ell t},0\}\\ +\delta V_{ij\ell}(\alpha R_{t-1}+1-\alpha,\tilde{p}_{\ell t})\end{array}\right)\\ +\sigma_0(R_{t-1},\tilde{p}_{\ell t})\underline{V}(\tilde{p}_{\ell t})\end{array}\right|\tilde{p}_{\ell t-1}\right].
\end{aligned}
$$

Here, $\underline{V}(\tilde{p}_{\ell t})$ and $\underline{U}(\tilde{p}_{\ell t})$ are the termination payoffs of the carrier and shipper, respectively, given the current spot rate $\tilde{p}_{\ell t}$.

Shippers' and carriers' expected payoffs from the repeated game form the basis for carriers' bidding and shippers' selection of the primary carrier in the auction. Two modeling assumptions help us discipline these expected payoffs. First, carriers bidding on the same lane differ only in the match-specific gains $(\psi_{ij\ell},\eta_{ij\ell})$ that they would generate in a relationship with the shipper. As a result, the shipper's and carrier's expected payoffs from a potential relationship 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. In contrast, 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. To highlight the role of per-transaction rents in the auction, we henceforth write the expected payoffs of a shipper and a carrier at the beginning of a relationship as functions of their rents: $U(R_0,\tilde{p}_{\ell 0}|\psi_{ij\ell}-p_{ij\ell},\eta_{ij\ell}+p_{ij\ell})$ and $V(R_0,\tilde{p}_{\ell 0}|\eta_{ij\ell}+p_{ij\ell})$, respectively.

Auction. For the auction stage, we focus on the class of *symmetric monotone equilibria*. The key advantage of this class is that it reduces 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. Denote by $\theta_{ij\ell}\equiv\psi_{ij\ell}+\eta_{ij\ell}\in\Theta$ the total match quality between shipper i and carrier j on lane ℓ, and denote by G_ℓ^θ its distribution. We define symmetric monotone equilibria formally as follows.

Definition 1. (Symmetric monotone equilibrium) In a symmetric monotone equilibrium, there exists a differentiable function $\mathbf{b}:\Theta\to\mathbb{R}$ such that:

(i) (Single indexing) The shipper's rent is $\psi_{ij\ell}-p_{ij\ell}=\mathbf{b}(\theta_{ij\ell})$. The carrier's rent is $\eta_{ij\ell}+p_{ij\ell}=\theta_{ij\ell}-\mathbf{b}(\theta_{ij\ell})$.

(ii) (Monotonicity) Both of these rents, $\mathbf{b}(\theta_{ij\ell})$ and $\theta_{ij\ell} - \mathbf{b}(\theta_{ij\ell})$, are strictly increasing in $\theta_{ij\ell}$.

(iii) (Optimal bidding) Carriers choose their effective bids optimally for all $\theta_{ij\ell}$,

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

(iv) (Optimal selection of the primary carrier) The shipper chooses the carrier j with the highest effective bid subject to the shipper's expected payoff from that relationship being no less than that of her outside option,

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

(v) (Optimal response to shipment requests) If j is the primary carrier, j uses the optimal strategy $\sigma_{ij\ell} : (R_{t-1}, \tilde{p}_{\ell t}) \mapsto \{\text{accept, spot, idle}\}$ in the repeated game, defined by equation (2) for carrier rent of $\theta_{ij\ell} - \mathbf{b}(\theta_{ij\ell})$.

Most substantive in this definition are the first two equilibrium selection conditions. Specifically, condition (i) says that in equilibrium, match-specific gains affect auction outcomes only via the total match quality. That is, the resulting rents of the shipper and the carrier depend on the competitiveness of the auction and on both sides' outside options, but not on how much each side contributes to the total match quality. As a result, total match quality is the sufficient statistic for match-specific gains in determining auction outcomes. In addition, condition (ii) says that in equilibrium, both the shipper and the carrier strictly benefit from higher total match quality. This allows the shipper to 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.[29]

Under conditions (i) and (ii), our bidding game can be transformed into an auxiliary first-price auction with symmetric bidding function $\mathbf{b} : \Theta \mapsto \mathbb{R}$. This transformation gives rise to condition (iii) for carriers' bidding and condition (iv) for shippers' selection of primary carriers. Condition (v) follows from our analysis of the repeated game. In Appendix B, we provide sufficient conditions for the existence of symmetric monotone equilibria and our empirical verification of these sufficient conditions.[30]

[29] Conditions (i) and (ii) 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.

[30] More precisely, the transformation of our bidding game into an auxiliary first-price auction on shipper

4.2 Market dynamics

While the individual relationship-level dynamics addressed in the previous subsection are central to identifying key model primitives—such as match-specific gains or search and operational costs—, our analysis of market-level welfare further depends on the dynamic interactions between relationships and the spot market as a whole. These market-level dynamics depend on (i) the stock of relationships and (ii) shocks to demand and capacity. In our welfare analysis, we keep (ii) fixed and vary the role and performance of relationships. This subsection explains how relationship dynamics evolve and how the market clears under the current market institution. The market dynamics under counterfactual institutions are detailed in Appendix G.

We treat as exogenous and fixed the following demand and capacity shocks: the measure of shippers who hold RFPs for long-term contracts on lane ℓ in period t, denoted $L^*_{\ell t}$; shocks to direct spot demand, denoted $D^*_{\ell t}$; and shocks to carrier capacity, denoted $C^*_{\ell t}$.[31] Conditional on these shocks, market dynamics are driven by two processes: first, the creation and termination of relationships and, second, carriers' decisions both within relationships and in the spot market. The former affects the *stock* of relationships and the latter affects the *states* of these relationships, summarized by the collection $\mathcal{H}_{\ell t-1} = (\psi_{ij\ell} - p_{ij\ell}, \eta_{ij\ell} + p_{ij\ell}, R_{ij\ell,t-1})_{ij\ell}$ of shippers and carriers' rents and carriers' rejection indices. We describe each of these processes in turn below.

First, the formation and termination of relationships determine the mass of relationships, the direct spot demand, and the direct spot capacity at the beginning of period t. We denote these objects respectively by $L_{\ell t}(\tilde{p}_{\ell t}|\mathcal{H}_{\ell t-1})$, $D_{\ell t}(\tilde{p}_{\ell t}|\mathcal{H}_{\ell t-1})$, and $C_{\ell t}(\tilde{p}_{\ell t}|\mathcal{H}_{\ell t-1})$. The change in the mass of relationships from one period to the next is determined by two factors: the fraction of shippers in $L^*_{\ell t-1}$ who successfully formed new relationships in period $t-1$ and the termination of existing relationships at the beginning of period t by shippers according to their punishment scheme σ_0. Recall that this punishment scheme depends on carriers' rejection indices (contained in $\mathcal{H}_{\ell t-1}$) and on the current spot rate $\tilde{p}_{\ell t}$.

The creation and termination of relationships also affect direct spot demand and capacity. On the demand side, shippers who failed to form new relationships in period $t-1$, along with those who terminated relationships at the beginning of period t, contribute to direct spot

rents also depends on the following information assumptions: $(\psi_{ij\ell}, \eta_{ij\ell})$ are drawn i.i.d. across carriers; $\psi_{ij\ell}$ is known by shipper i and carrier j; $\eta_{ij\ell}$ is privately known by carrier j. These assumptions imply 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, $\mathbf{b}(\theta_{ij\ell}) = \psi_{ij\ell} - p_{ij\ell}$.

[31] Our model does not take a stance on why shippers might want to hold RFPs versus not, since we do not observe this decision margin. In practice, shippers who hold RFPs tend to be larger shippers with relatively consistent demand for shipments over a long period of time. Similarly, we treat shocks to trucking capacity as exogenous—an assumption that is likely to hold in the medium (but not long) run.

demand in period t. On the supply side, carriers who formed new relationships exit the pool of spot carriers, while those whose relationships were terminated re-enter this pool. Appendix G details this evolution of demand and capacity over time under additional assumptions on the evolution of the underlying shocks $(L^*_{\ell t}, D^*_{\ell t}, C^*_{\ell t})$.[32]

Second, carriers' decisions determine realized relationship volume and spot volume, ultimately clearing the market. Recall from equation (2) that within a relationship, the carrier's decision to accept the shipper's load, reject it to service the spot market, or reject it to remain idle depends on which is the highest among the following: the carrier's full compensation $\bar{p}$ for an acceptance, the spot rate $\tilde{p}_{\ell t}$, or the sum $c_{j\ell t}+\kappa_\ell$ of operational and search costs. Thus, the aggregate share of carriers within relationships that accept relationship loads depends on the distribution $\mu(\cdot|\tilde{p}_{\ell t}, \mathcal{H}_{\ell t-1})$ of full compensation $\bar{p}$ across relationships. Additionally, carriers without relationships service spot loads only if their costs are sufficiently low. Together, these result in the following spot-market clearing condition:

$$\begin{aligned} &\underbrace{L_{\ell t}(\tilde{p}_{\ell t} \mid \mathcal{H}_{\ell t-1}) \left(1 - \int_{\tilde{p}_{\ell t}}^{\infty} F_\ell(\bar{p} - \kappa_\ell) d\mu_{\ell t}(\bar{p}|\tilde{p}_{\ell t}, \mathcal{H}_{\ell t-1})\right)}_{\text{Overflow demand}} + \underbrace{D_{\ell t}(\tilde{p}_{\ell t} \mid \mathcal{H}_{\ell t-1})}_{\text{Direct spot demand}} \\ = \quad &(\underbrace{L_{\ell t}(\tilde{p}_{\ell t} \mid \mathcal{H}_{\ell t-1}) \mu_{\ell t}(\tilde{p}_{\ell t}|\tilde{p}_{\ell t}, \mathcal{H}_{\ell t-1})}_{\text{Overflow capacity}} + \underbrace{C_{\ell t}(\tilde{p}_{\ell t} \mid \mathcal{H}_{\ell t-1})}_{\text{Direct spot capacity}}) F_\ell(\tilde{p}_{\ell t} - \kappa_\ell). \end{aligned} \tag{4}$$

This equation captures the crowding-out effect of the spot market on relationships: carriers whose full compensations are lower than the current spot rate go to the spot market, increasing spot capacity, while shippers' loads that are rejected within relationships must be fulfilled in the spot market.

Finally, to discipline the crowding-out effect of relationships on the spot market, we assume that search cost $\kappa_\ell = \kappa(\text{Volume}_\ell^{\text{spot}})$ is a function of the average spot volume. This scale-efficiency relation, combined with the market clearing conditions in each period and the relationship dynamics described above, determines the equilibrium search cost κ_ℓ and the equilibrium path $\{\tilde{p}_{\ell t}\}$ of spot rates.[33]

[32] These assumptions are imposed for the estimation of the underlying demand and capacity shocks, as described in Appendix G.

[33] Note that depending on the dynamics of the underlying demand and capacity shocks $(L^*_{\ell t}, D^*_{\ell t}, C^*_{\ell t})$, the equilibrium path of spot rates can often be reasonably approximated by an autocorrelated process. The assumption that individual shippers and carriers perceive the spot rate process as AR(1), rather than a higher-order autocorrelated process, is made to simplify estimation by limiting the state space in carriers' discrete choice problem. Similarly, search costs could, in principle, depend on the current spot volume $\text{Volume}_{\ell t}^{\text{spot}}$ rather than its average over time. However, allowing for this would require expanding the state space to include current spot volume. This would pose no problems for identification, but would complicate estimation.

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.[34] 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 $G_\ell^{\psi,\eta}$ of match-specific gains across relationships. Our identification argument relies on the following assumptions, which will be maintained throughout our analysis.

Assumption 2. *(Regularity) Assume the following regularity conditions:*

(i) The AR(1) spot process $\mathcal{P}_\ell$ as perceived by individual shippers and carriers has support $\mathbb{R}^+$ for all $\tilde{p}_{\ell t-1}$.

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

(iii) The underlying distribution of match-specific gains $G_\ell^{\psi,\eta}$ has full support in $\mathbb{R}^2$. Moreover, it induces an underlying distribution of match quality $G_\ell^\theta \equiv G_\ell^{\psi+\eta}$ that has a strictly decreasing hazard rate, $g_\ell^\theta / G_\ell^\theta$.[35]

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

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

Assumption 3. *(Properties of full compensations) Given the perceived spot process $\mathcal{P}_\ell$ and the punishment scheme σ_0, the full compensation $\bar{p}(R_{t-1}, \tilde{p}_{\ell t} | \eta + p)$ is:*

(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 3 is essentially an assumption on the underlying spot process and the shipper's punishment scheme. The most substantive assumption is 3(i), which can be numerically verified. Assumption 3(ii) holds under mild regularity conditions, and Assumption 3(iii) is satisfied if the shipper's punishment scheme σ_0 is strictly increasing in the carrier's rejection index R_{t-1}.

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

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

Theorem 1. *(Full identification) Under Assumptions 2 and 3 and within the class of symmetric monotone equilibria, $(\sigma_0, N, (\mathcal{P}_\ell, G_\ell^{\psi,\eta})_\ell)$ are identified, and cost parameters $(F_\ell, \kappa_\ell)_\ell$ are identified up to a constant.*

The perceived spot process $\mathcal{P}_\ell$ is identified as the AR(1) process that best fits the realized path of spot rates. The shipper's punishment scheme σ_0 is identified, since every Markov state $(R_{t-1}, \tilde{p}_{\ell t})$ is observed under Assumption 2.[36] The number N of bidders is identified from the number of carriers observed in the routing guide. Given these instrumental objects, we now identify the distribution F_ℓ of operational costs, search cost κ_ℓ, and the joint distribution $G_\ell^{\psi,\eta}$ of shippers and carriers' match-specific gains on each lane. In the next subsections, we describe the key lemmas and delegate the formal proof of Theorem 1 to Appendix A.

5.1 Identifying 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 $\tilde{c}_{j\ell t} = c_{j\ell t} + \kappa_\ell \sim \text{Normal}(\mu_\ell^c + \kappa_\ell, \sigma^c)$.[37] 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 punishment scheme 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 $(R_{t-1}, \tilde{p}_{\ell t})$,

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

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 (F_ℓ, κ_ℓ). 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

[36]Specifically, in Assumption 2, (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.

[37]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

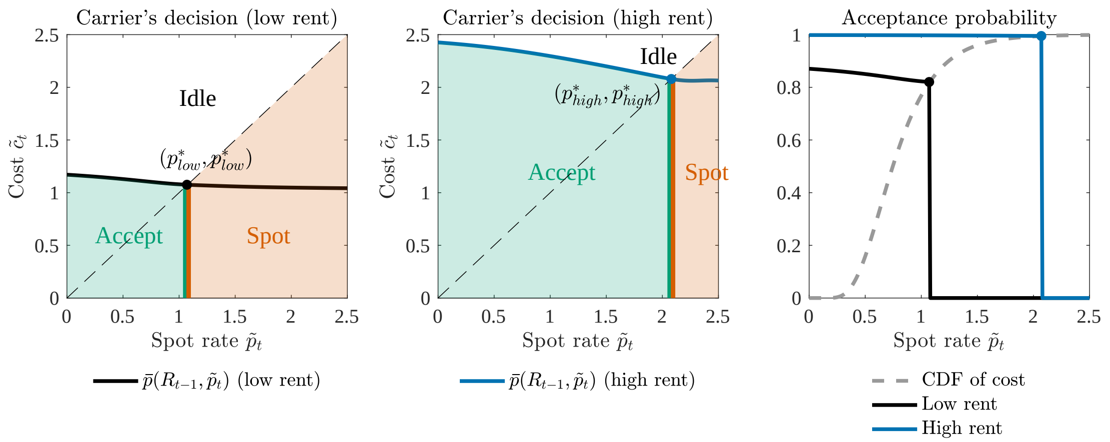

rate (x axis) and transformed cost (y axis). As spot rate increases, the gap between the full compensation associated with each rent level and the spot rate shrinks, with the full compensation schedule crossing the 45-degree line at a critical point p^*. The carrier decides between *accept* and *idle* 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 at p^*, the acceptance probability—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,

$$\lim_{\tilde{p}_t \uparrow p^*} \Pr(d_t = \text{accept}|R_{t-1}, \tilde{p}_t) = \tilde{F}_\ell(p^*) = \lim_{\tilde{p}_t \downarrow p^*} \Pr(d_t = \text{spot}|R_{t-1}, \tilde{p}_t) > 0. \tag{5}$$

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, $p^*_{low} < p^*_{high}$.[38] In summary, within-lane variation in carrier rent 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.[39] To overcome this

[38]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.

[39]On-path terminations of relationships occur due to demotion and auction events, resulting in short-lived

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.[40] 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 transformed costs into search costs and 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.

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 $\text{Volume}_\ell^{\text{spot}}$ identifies search costs $\kappa_\ell = \kappa(\text{Volume}_\ell^{\text{spot}})$ and the distribution F_ℓ of operational costs, up to a constant.*[41]

5.2 Identifying shipper primitives from the auction

We complete our identification argument with the identification of the joint distribution of match-specific gains for shippers and carriers in realized relationships. Crucially, we

relationships. For long-lasting relationships, we can identify the carrier's acceptance tendency conditional on the relationship surviving to the next period. Appendix F provides details on how we correct for this selection.

[40] 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.

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

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, Perrigne and Vuong (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.

With $\mathbf{b}_r$, we pin down the distribution of shipper rent from the previously identified distribution of carrier rent. Conditioning the joint distribution of shipper and carrier rents on observed 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 is identified.*

6 Estimation

This section outlines our estimation procedure and presents 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 F explains our estimation procedure in greater 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 rather the time-specific mean $\tilde{p}_{\ell t}$ and standard deviation σ_ℓ^ζ, we assume that at the time of decision-making, carrier j faces spot rate $\tilde{p}_{\ell t}+\zeta_{j\ell t}$, with $\zeta_{j\ell t} \sim \text{Normal}(0, \sigma_\ell^\zeta)$. We estimate an AR(1) process for the calendar-based spot rate process and scale it by the frequency of shipper-carrier interactions in a relationship to obtain the load-based spot process in that relationship. We similarly scale the daily discount rate of 0.992 to obtain the relationship-specific, load-based, discount rate δ_{ial}.

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

$$\sigma_0(R_{t-1}, \tilde{p}_{\ell t}) = \Phi(\alpha_1 + \alpha_2 R_{t-1} + \alpha_3 \mathbf{X}_{ia\ell} + \alpha_4 R_{t-1}\mathbf{X}_{ia\ell} + \alpha_5(\tilde{p}_{\ell t} - \tilde{p}_{\ell 0})), \tag{6}$$

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.[42] 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 (2025), 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 E.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

Estimation procedure. We estimate cost parameters under the following parametric specification. Assume that transformed costs $\tilde{c}_{j\ell t} \sim \text{Normal}(\tilde{\mu}^c_{ia\ell}, \sigma^c)$ with auction-specific

[42] As argued in Harris and Nguyen (2025), the frequency and timing consistency of load offers, as well as changes in spot market conditions, influence the continuation value of a relationship. Thus, these factors may affect the harshness of the shipper's punishment in response to carrier rejections.

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:

$$\begin{aligned}\tilde{\mu}^c_{ia\ell} = &\overbrace{\gamma_0 + \gamma_1 \ln(\text{Volume}^{\text{spot}}_\ell) + \gamma_2 \ln(\text{Volume}^{\text{spot}}_\ell) \ln(\text{Distance}_\ell)}^{\text{search cost } \kappa_\ell} \\ &+ \underbrace{\gamma_3 \ln(\text{Distance}_\ell) + \gamma_4 \mathbf{X}_{ia\ell} + \gamma_5 \text{Tight}_a + \gamma_6 \text{Imbalance}_\ell + \nu^c_\ell + \epsilon^c_{ia\ell}}_{\text{mean operational cost } \mu^c_{ia\ell}}.\end{aligned} \tag{7}$$

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.[43]

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.[44] 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 (7) 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).[45] 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 2 presents our estimates of cost parameters. Figure 4 plots the estimated search costs and mean of operational costs

[43] In principle, search frictions might vary with market phases (e.g., Vreugdenhil, 2023). Our specification subsumes such effect under the reduced-form term Tight_a. This choice is due to our demand shifter—or our instrument for spot volume—having only cross-sectional variation and not temporal variation. Another concern is that 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. In our setting, including $\text{Imbalance}_\ell \equiv \ln(\text{Volume}^{\text{spot}}_{-\ell}) - \ln(\text{Volume}^{\text{spot}}_\ell)$ in equation (7) helps mitigate such spatial-equilibrium 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.

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

[45] 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, which relies mostly on our estimates of the scale efficiency parameters.

Table 2: 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	$(0.7, 1.1)$

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

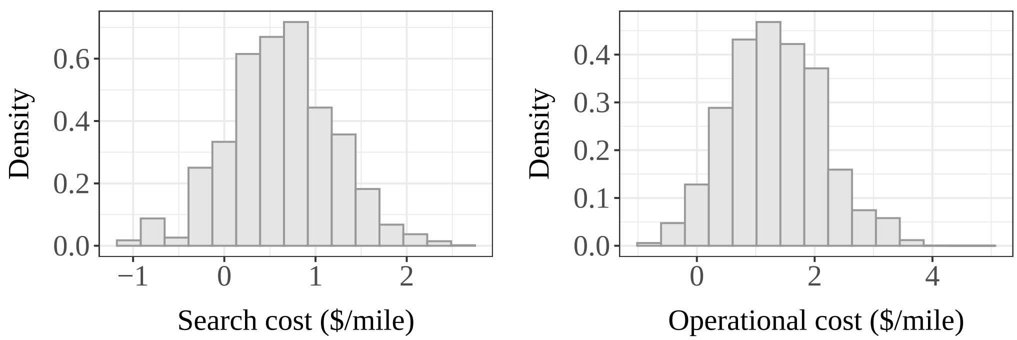

across all relationships.

Large spot market scale efficiency. The estimates in Table 2 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.

Determinants of operational costs. Table 2 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 for which finding a backhaul is twice as easy. These factors contribute to the large variation in mean operational costs across relationships in Figure 4.

6.3 Relationships' match-specific gains

Estimation procedure. Given our cost estimates, we recover match-specific gains non-parametrically 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 $\tilde{\mathbf{X}}_{ia\ell}$ of observable characteristics that includes auction-specific characteristics in $\mathbf{X}_{ia\ell}$ and the following lane-specific characteristics: Rate_ℓ, $\text{Volume}_\ell^{\text{spot}}$, Distance_ℓ, Tight_a, and Imbalance_ℓ.[46] 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, Perrigne and Vuong (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 $\mathbf{b}_r(\underline{r}_{ia\ell}) = -\underline{r}_{ia\ell} - \kappa_\ell$.[47] 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. 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: even if the spot market could be made so thick as to completely eliminate search costs, four-fifths of current relationships would still have value. These benefits, however, are very heterogeneous across relationships, a fact that underscores the importance of auctions in facilitating the formation

[46] The idea is to use observed lane-specific characteristics, including equilibrium objects such as spot rate and volume, to control for differences in the underlying demand and supply factors across lanes.

[47] 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

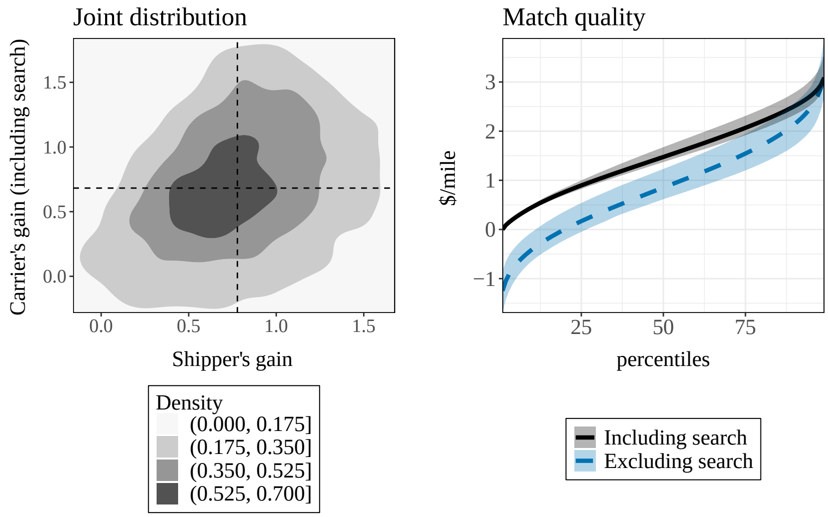

Note: The bands on the right panel are 95% bootstrap confidence intervals.

of high-value matches in the truckload setting.

7 Welfare analysis

In this section, we combine our estimates of the model primitives from the previous section with a market equilibrium framework to answer the following question: Should relationships play a more or less prominent role in the US truckload freight industry? Our findings suggest that a marginal—but not global—shift away from relationships would improve social welfare.

7.1 Conceptual framework

Our welfare analyses focus on two frictions in the current market institution: (i) the contractual frictions within long-term relationships and (ii) the thinness—and therefore inefficiency—of the spot market.

To assess the role of contractual frictions, we evaluate the welfare implications of replacing the current fixed-rate contracts with index-priced contracts that are designed to maximize welfare at the level of individual relationships. These contracts build on the standard idea in contract theory (e.g. Harris and Raviv, 1979; Shavell, 1979) that, to solve a moral hazard problem in a principal-agent relationship, the principal should "sell the firm" to the agent. Specifically, these index-priced contracts transfer all relationship rents to the carrier, peg payments one-to-one to the spot rate, and award to the carrier that bids the highest fixed fee a contract with the following payment schedule:[48]

$$p_{ij\ell}(\tilde{p}_{\ell t}) = \begin{cases} -\,b^0_{ij\ell} + \overbrace{\psi_{ij\ell} + \tilde{p}_{\ell t}}^{\text{“incentive”}} & \text{if carrier } j \text{ accepts} \\ \underbrace{-b^0_{ij\ell}}_{\text{“screening”}} & \text{if carrier } j \text{ rejects.} \end{cases} \tag{8}$$

While these indexed-price contracts eliminate contractual frictions within relationships, their widespread adoption would further thin the spot market, thereby reducing its efficiency. Comparing the market-level performance of these index-priced contracts to that of fixed-rate contracts helps shed light on whether market inefficiency in the current institution stems primarily from contractual frictions or primarily from the lack of spot market scale.

More systematically, we assess the key market-level tradeoff—between realizing match-specific gains within relationships and maintaining spot market efficiency—across a set of counterfactuals in which we eliminate contractual frictions while varying the split between relationship and spot transactions. Such counterfactuals not only highlight the fundamental tradeoff between relationships and the spot market, but are also particularly tractable.[49] For a given level of spot market share in total market volume, we can find the optimal allocation between relationship and spot transactions by solving a linear programming problem. As shown in Appendix G, the solution to this constrained first-best problem is implementable by pairing index-priced contracts with an appropriate tax on relationship transactions. Intuitively, index-priced contracts address the externalities of carriers' decisions on shippers within relationships, while varying the tax rate mitigates the externalities of relationships on the spot market.

Comparing the current institution to the constrained first-best curve—which traces the highest achievable welfare for each level of spot share—enables us to determine whether the market-level first-best relies more or less on relationships relative to the current institution. In addition to the market-level first-best, we also zoom in on two other benchmarks: (i) an *index-priced* scenario without taxes, as previously mentioned, and (ii) a *spot-only* scenario, which is equivalent to imposing an infinite tax on relationship transactions. These benchmarks represent two practically relevant cases: one where practitioners shift toward

[48]Under this 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.

[49]By removing contractual frictions—and thus the need for relational incentives—our counterfactuals significantly simplify relationship dynamics. Additionally, since these counterfactuals do not rely on relational incentives, we bypass the need to explicitly model shippers' rationale for their observed punishment schemes.

more flexible contracts and another where an Uber-like model with a centralized spot market dominates trucking.

To ensure a meaningful comparison across these scenarios, we estimate and hold constant the underlying aggregate demand and capacity factors. To do so, we make the following assumptions. First, KMA-to-KMA lanes can be grouped into seven lane types.[50] Second, within each lane type, 2016 represents a soft market, and 2018 represents a tight market, resulting in a total of fourteen lane-year clusters. Third, within each of these clusters, the underlying demand and capacity are characterized by the tuple $(L^*, C^*, (D_t^*)_{t=1}^{52})$. We assume that all L^* relationships are formed at $t = 0$ and that total carrier capacity remains constant at $L^* + C^*$ throughout the year. Additionally, D_t^* represents an exogenous shock to spot demand in week t, capturing variation in spot rates not explained by observed relationship dynamics. Fourth, we assume that loads fulfilled by backup carriers can be treated as spot loads, based on the observation that, compared to primary carriers, backup carriers are substantially less likely to accept offers and perform substantially worse on quality metrics.[51] We estimate the tuple $(L^*, C^*, (D_t^*)_{t=1}^{52})$ of underlying demand and capacity using our model estimates in Section 6 along with observed rates and volumes in relationship and spot transactions. Appendix G provides a detailed explanation of this estimation process and the derivation of welfare in both the current and counterfactual scenarios.

7.2 Counterfactual welfare

Index-priced (without taxes). We assess the role of contractual frictions at both the relationship and market levels by quantifying (i) the effect of index-priced contracts on a single relationship and (ii) the market-level welfare effects if all relationships move to index-priced contracts. At the relationship level, we find that current relationships generate an average surplus of \$0.64/mile over spot transactions. While large relative to average spot and contract rates, this surplus amounts to only 64% of the first-best surplus achieved under index-priced contracts. Moreover, the share of first-best surplus achieved is significantly lower in relationships with lower match quality.[52]

[50] 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 cluster lanes based on exogenous and equilibrium characteristics. Figure 14 in Appendix G visualizes the effectiveness of our approach in separating lanes with different underlying supply and demand factors.

[51] In Appendix E.1, we show that backup asset-owners have acceptance rate roughly half that of their primary counterparts. Moreover, their performance on on-time delivery and detention time is quite similar to that of brokers, who match shippers with different carriers—potentially including spot carriers—on a per-load basis.

[52] This selection effect—driven by higher-quality relationships facing a less severe moral hazard problem and having greater flexibility to use contract rates as an incentive instrument—softens the impact of contractual frictions on aggregate welfare. In Appendix G, we provide more details on the welfare effects of fixed-rate

Figure 6: Welfare change relative to the current institution

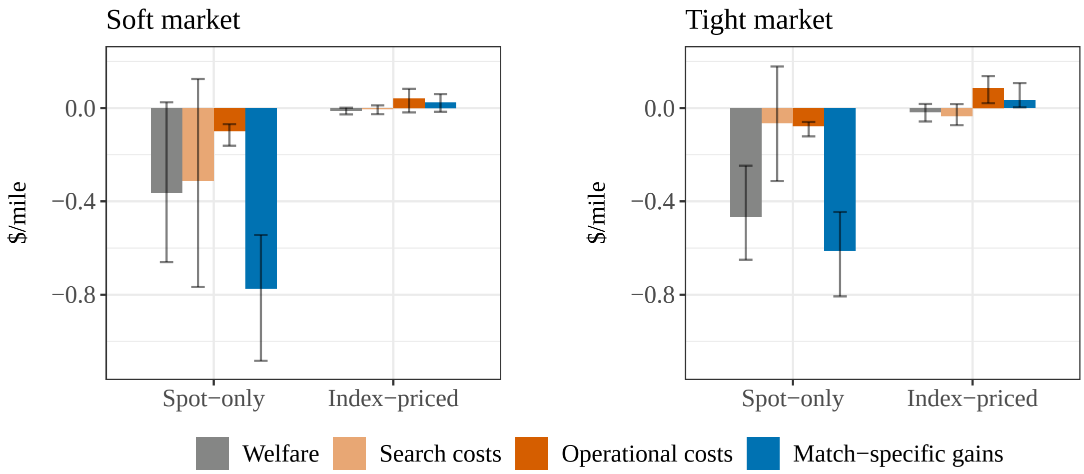

Note: 95% confidence intervals are constructed from 50 full-estimation bootstraps with auction-level sampling.

At the market level, the chief drawback of index-priced contracts is their tendency to further thin the spot market and increase search costs for spot loads. We find that the global adoption of index-priced contracts would increase search costs by \$0.09/mile in a soft market and \$0.22/mile in a tight market. These higher search costs affect welfare through three channels: a direct channel (search costs), and two indirect channels (operational costs and match-specific gains).

Figure 6 presents welfare changes (\$/mile) of the *index-priced* scenario relative to the current institution across these channels. For the direct channel, we find that total search costs incurred remain virtually unchanged. While the cost of each search is higher, fewer carriers participate in the spot market, meaning that fewer carriers incur these costs. For the indirect channels, higher search costs increase the effective premium of relationships relative to the spot market. This in turn leads to (i) the formation of more relationships with negative intrinsic match quality and (ii) the acceptance of loads within relationships by carriers with a higher operational cost draw than some carrier operating in the spot market. The first of these indirect effects dampens the gains from increasing transactions with positive intrinsic match quality, while the second increases average operational costs. Both indirect effects tend to offset the gains from resolving contractual frictions by index-priced contracts.

In sum, our findings on the welfare effects of index-priced contracts suggest strong private incentives to adopt more flexible contracts while cautioning against their potential negative

versus index-priced contracts for individual relationships, examining (i) the impact of match quality, (ii) the influence of market conditions, and (iii) the distribution of relationship surplus between shippers and carriers.

impact on spot market thickness and efficiency.[53]

Spot-only. Turning from a scenario in which relationships are optimized to one in which they are eliminated, we find that the *spot-only* scenario would result in a substantial welfare loss relative to the current institution (see Figure 6). Specifically, this loss amounts to $0.36/mile in a soft market and $0.47/mile in a tight market. While the reductions in search and operational costs are large, they fall well short of compensating for the complete loss of match-specific gains from relationships. This finding confirms the importance of relationship benefits in trucking. A global shift to a dominant spot platform is not only unlikely—due to private incentives to remain in relationships—, but also socially inefficient.

Market-level first-best. Finally, we compare the current institution to the market-level first-best. The left panel of Figure 7 illustrates the constrained first-best curve for a cluster of lanes in the soft market phase and maps different institutions along this curve.[54] Three patterns emerge, all of which similarly hold across other clusters of lanes and market phases.[55] First and most critically, comparing the *fixed-rate* and first-best points reveals that the current institution features a suboptimal spot share; using the tax instrument to marginally increase the spot share would improve welfare. Second, implementing *index-priced* contracts (with no tax) would further decrease the spot share, moving it further away from the first-best level. This suggests that the incompleteness of current contracts, by limiting relationship surplus, mitigates the negative externalities relationships exert on the spot market. Third, while the first-best spot share is higher than that of the current institution, relationships continue to play a key role in the first-best. Indeed, *spot-only* is the worst-performing market institution along the constrained first-best curve. This underscores the indispensable role of relationships in the trucking industry.

The right panel of Figure 7 shows how far the current institution is from the market-

[53] Conversations with practitioners suggest that, historically, the lack of access to reliable freight indices with granular pricing and location data has been a barrier to implementing index-priced contracts. In recent years, there have been early signs of interest in index-priced contracts, with the introduction of those that utilize a national truckload index. See, for example, https://www.freightwaves.com/news/index-linked-contracts-new-way-to-buy-trucking-capacity.

[54] Technically, the current institution with fixed-rate contracts is not allocatively efficient. However, its proximity to the constrained first-best curve suggests that most inefficiencies of the current institution stem from the suboptimal split between relationship and spot transactions.

[55] The finding that the *index-priced* scenario results in suboptimal spot share and that the *spot-only* scenario is the worst-performing institution holds across all 14 clusters. The finding that the *fixed-rate* scenario falls between the *index-priced* scenario and the market-level first-best—thus outperforming the *index-priced* scenario in market-level welfare—holds in 12 out of 14 clusters. Another general pattern is that the portion of the constrained first-best curve between *index-priced* and the market-level first-best is relatively flat, suggesting that achieving the first-best would require a large tax while delivering a relatively small welfare gain. Across our 14 clusters, the first-best tax ranges from $0.38/mile to $0.54/mile.

Figure 7: Comparison to market-level first-best welfare

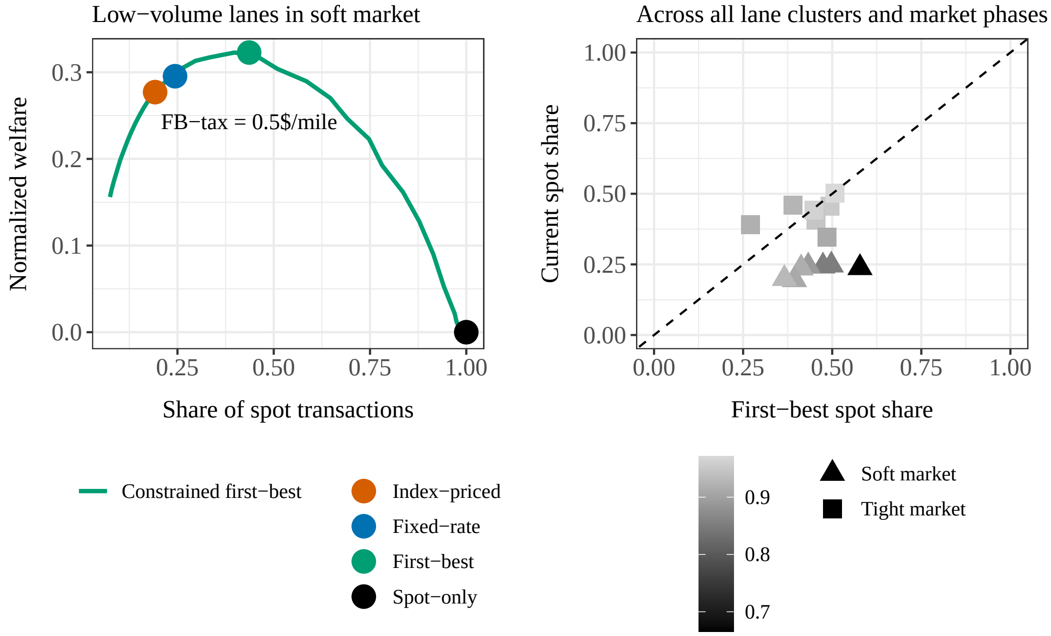

Note: The color bar on the right panel shows the fraction of the market-level first-best surplus that is achieved by the current institution. For meaningful comparisons across clusters of lanes and market phases, we normalize market-level surplus by the market-level welfare achieved under the *spot-only* scenario, which is universally worst for social welfare across all clusters of lanes and market phases.

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. 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 better in a tight market, when pressure from the spot market is high. That is, the incompleteness of fixed-rate contracts acts as a partial corrective tax on relationship transactions, a mechanism that works particularly well in high-demand periods.

7.3 Discussion

Reflecting on our empirical framework and on the results presented above, we offer four remarks. The first two remarks are conceptual and provide additional nuance to the interpretation of our empirical framework. The third and fourth remarks are empirical, highlighting additional conclusions that can be drawn from our estimated framework. We expound upon these additional empirical conclusions in Appendix E.3.

First, while we study two-way crowding-out effects between relationships and the spot market as hypothesized by Kranton (1996), we emphasize a slightly different problem. In

Kranton (1996), such crowding-out effects can lead to multiple equilibria, potentially resulting in either too much or too little relationship activity—a *coordination* failure in equilibrium selection. In contrast, our counterfactual framework emphasizes a *cooperation* problem. We consider a social planner who uses external instruments, such as taxes, to alter the equilibrium balance between relationship and spot transactions.[56]

Second, while we refer to the spot market scale (in)efficiencies as "search costs," we want to emphasize that we allow for a much broader interpretation of the estimated scale efficiencies. Indeed, $\kappa(\text{Volume}_\ell^{\text{spot}})$ can be thought of as a reduced form relationship capturing any number of mechanisms by which spot market participants may benefit from greater market thickness. To name just a few possible mechanisms, a thicker market may facilitate faster matching or increase the number of draws from the match-quality distribution, thereby increasing the expected maximum match-quality drawn.[57] We are agnostic as to which of these microfoundations gives rise to our estimated scale efficiencies.

Third, while this section focuses on comparisons of aggregate welfare across institutions, this aggregation masks significant distributional effects. Increasing spot share improves total welfare, and its biggest winners are spot carriers, typically small asset owners, who benefit from increased demand for spot loads and lower search costs. The biggest losers, on the other hand, are shippers who prefer relationships, as they both forgo match-specific gains and have reduced bargaining power in contract formation because carriers enjoy improved outside options.

Finally, while our counterfactual framework holds the extensive margin fixed to focus on the balance between relationships and the spot market, the direction of price effects enables us to determine the potential welfare gains with elastic demand for trucking services. We estimate that a higher spot share lowers equilibrium spot rates due to reduced search costs and improved allocative cost efficiency. This suggests that a higher spot share would increase overall trucking volume were demand to be elastic. Appendix E.3 provides a more detailed discussion of these price and distributional effects, focusing on the comparison with the

[56]Our counterfactual approach—benchmarking current welfare against the constrained first-best welfare frontier—readily generalizes to other markets that are characterized by relationship gains and spot market scale efficiencies. Constructing this frontier involves two steps: (i) deriving the first-best arrangements within relationships, and (ii) solving for the relationship tax that implements each level of spot market share. By tracing the optimal allocation at each relationship-spot split, this approach sidesteps equilibrium selection issues that arise from the endogenous link between spot market thickness and efficiency. As in our setting, comparing the current institution to this frontier reveals the extent of frictions in existing relationships and whether relationships should play a more or less prominent role in the overall market.

[57]In addition to improving expected match quality conditional on the average quality of spot shippers/-carriers, a similar effect may arise as a result of positive selection. For example, if—as seems likely—the highest quality shippers/carriers are more likely to be in relationships, anything that decreases relationship premiums (e.g., a decrease in search costs) would likely attract to the spot market participants whose quality exceeds the current average spot-market quality.

index-priced and *spot-only* scenarios.

8 Conclusion

This paper studies the market-level tradeoff between long-term relationships and the spot market, asking whether the trucking industry—in which long-term relationships are currently dominant—would benefit from a shift toward more spot transactions. To answer this question, we develop an empirical framework to quantify the intrinsic gains from relationships and their externalities on the spot market. This framework captures the critical tension between relationships and spot, similar to the two-way crowding-out effects hypothesized by Kranton (1996). On the one hand, the current relationships feature contractual frictions, thus allowing spot temptation to crowd out relationships. On the other hand, spot markets feature frictions that are worsened the more transactions occur within relationships rather than in the spot market.

Using this framework, together with detailed data on the US for-hire truckload freight industry, we quantify these forces and find that relationships exert a strong negative scale efficiency externality on the spot market. Locally, this externality dominates, so a *marginal* shift toward spot transactions would increase social welfare. However, we also find that relationships capitalize on large intrinsic gains. Thus, a *global* shift toward spot in which all relationships were abandoned in favor of spot transactions would result in large welfare losses.

Our findings also highlight the potential benefits of technological advances—such as algorithmic matching and pricing—that improve spot market efficiency. While these advances would clearly deliver direct benefits to spot market participants, they would also have the indirect effect of shifting the equilibrium split away from relationships and toward spot transactions. Ex ante, the welfare effects of this shift are ambiguous, since relationship gains may be lost, but spot market efficiency may be improved. In the case of for-hire truckload freight, however, our results indicate that this is a shift toward the market-level first-best. In other words, these indirect effects would amplify the direct benefits of technological advances in the spot market for truckload freight.

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 2 and 3 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 3(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 3(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. 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 $(\tilde{\mu}^c,\sigma^c)$. □

Proof of Lemma 2. 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 $\lim_{r\uparrow\bar{r}}[g^{\eta+p}]^{1:N}(r)>\lim_{r\uparrow\bar{r}}[\hat{g}^{\eta+p}]^{1:N}(r)$. 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

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

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. 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 $r = \mathbf{r}(\theta) \geq \underline{r}$,

$$(N-1)\frac{g^{\eta+p}(r)}{G^{\eta+p}(r)} = \frac{\frac{\partial}{\partial r}V(R_0, \tilde{p}_0|r)}{V(R_0, \tilde{p}_0|r) - \mathbf{E}_{F,\mathcal{P}}[\underline{V}(\tilde{p}_1)|\tilde{p}_0]}\mathbf{b}'_r(r). \tag{9}$$

Notice that the carrier's expected payoff as a function of carrier's rent r is identified from the punishment scheme σ_0, cost distribution F and search cost κ. To pin down $\mathbf{b}'_r$ from equation (9), 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

$$[G^{\eta+p}]^{1:N}(r) = \frac{[G^{\eta+p}(r)]^N - [G^{\eta+p}(\underline{r})]^N}{1 - [G^{\eta+p}(\underline{r})]^N}. \tag{10}$$

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 (10).

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

$$U(R_0, \tilde{p}_0|\mathbf{b}(\underline{r}), \underline{r}) = \mathbf{E}_{\mathcal{P}}[\underline{U}(\tilde{p}_1)|\tilde{p}_0]. \tag{11}$$

Note that $\underline{r}$ is identified as the lowest point in the support of $G_r^{\eta+p}$, and the RHS of equation (11) is identified from the spot process. Thus, equation (11), with the LHS being strictly increasing in the shipper's rent, pins down $\mathbf{b}_r(\underline{r})$. □

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.

Acocella, Angela, Chris Caplice, and Yossi Sheffi. 2022*a*. "The end of set-it-and-forget-it pricing? Index-based freight contract design under uncertainty." MIT Center for Transportation & Logistics Working Paper Series.

Acocella, Angela, Chris Caplice, and Yossi Sheffi. 2022*b*. "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." National Bureau of Economic Research Working Paper.

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.

Asker, John, and Estelle Cantillon. 2008. "Properties of scoring auctions." *The RAND Journal of Economics*, 39(1): 69–85.

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.

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

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.

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

Berdinis, Eric. 2024. "Uber Freight Exchange: Envisioning a freight marketplace for all." Technical Report, Uber Freight.

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

BLS. 2024. "Contingent and Alternative Employment Arrangements News Release." Technical Report, U.S. Bureau of Labor Statistics.

Bokher, Sam. 2018. "Segments of U.S. trucking." Technical Report.

Bonhomme, Stéphane, and Elena Manresa. 2015. "Grouped patterns of heterogeneity in panel data." *Econometrica*, 83(3): 1147–1184.

Boudreau, Laura, Julia Cajal-Grossi, and Rocco Macchiavello. 2023. "Global value chains in developing countries: A relational perspective from coffee and garments." *Journal of Economic Perspectives*, 37(3): 59–86.

Brancaccio, Giulia, Myrto Kalouptsidi, and Theodore Papageorgiou. 2020. "Geography, transportation, and endogenous trade costs." *Econometrica*, 88(2): 657–691.

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." Brown University Working Paper.

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.

Castillo, Juan Camilo. 2023. "Who benefits from surge pricing?" Working Paper.

CHR. 2015-2019. "TMC Tendering Data: 2015-2019." *C.H. Robinson.*

DAT. 2015-2019. "Lane-Level Truckload Dry Van Spot Rate Data: 2015-2019." *DAT Solutions.*

Faltings, Richard. 2025. "Pricing the Right to Renege in Search Markets: Evidence from Trucking." Working paper.

FMCSA. 2021. "MCMIS Census of Carriers." *Federal Motor Carrier Safety Administration.*

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.

Fudenberg, Drew, and Eric Maskin. 1986. "The folk theorem in repeated games with discounting or with incomplete information." *Econometrica*, 54(3): 533–554.

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.

Garcia-Osipenko, Maria, Nicholas Vreugdenhil, and Nahim Bin Zahur. 2025. "Endogenous Rigidities and Capital Misallocation: Evidence from Containerships." Arizona State University Working Paper.

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

Gavazza, Alessandro. 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. 2025. "Long-term relationships in the US truckload freight industry." *American Economic Journal: Microeconomics*, 17(1): 308–353.

Harris, Milton, and Artur Raviv. 1979. "Optimal incentive contracts with imperfect information." *Journal of economic theory*, 20(2): 231–259.

Holm, Seth. 2020*a*. "How big is the contract market?" Technical Report.

Holm, Seth. 2020*b*. "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.

Macchiavello, Rocco, and Ameet Morjaria. 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.

NMFTA. 2021. "Standard Carrier Alpha Codes." *National Motor Frieght Traffic Association.*

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.

Rose, Nancy L. 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.

Shavell, Steven. 1979. "Risk sharing and incentives in the principal and agent relationship." *The Bell Journal of Economics*, 55–73.

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.

Uber Technologies, Inc. 2025. "Uber Announces Results for Fourth Quarter and Full Year 2024." Technical Report.

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, American Transportation Research Institute.

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.

Supplemental Appendix

B Other proofs

B.1 Monotonicity in rents

Assumption 4. *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 holding her rent fixed, 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 8 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 3(i)). The right panel shows that the shipper's expected payoff is increasing in both her rent and the carrier's rent (Assumption 4).

B.2 Existence of a symmetric monotone equilibrium

Lemma 7. *Under Assumptions 2-4, 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. 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 4. 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

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

Figure 8: Monotonicity of carriers' full compensation and shippers' expected payoffs in rents

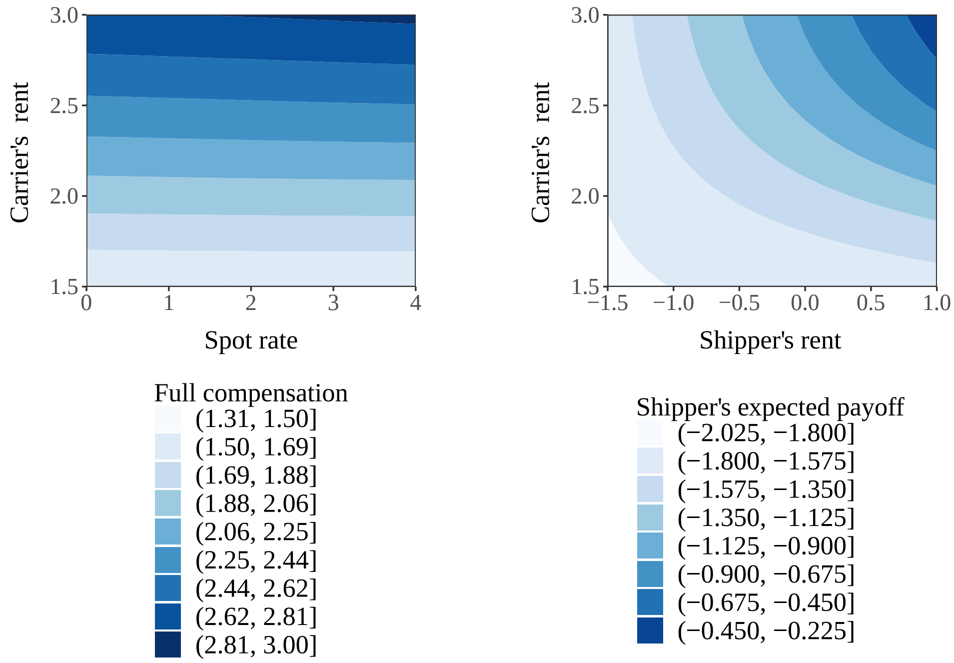

Note that a relationship strictly benefits the carrier if and only if the carrier's rent is strictly positive. Thus, in this equilibrium, the lowest match quality of a winning carrier gives zero rent to that carrier, $\mathbf{b}(\underline{\theta}) = \underline{\theta}$. That is, individual rationality 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. 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 $\mathbf{r}'(\underline{\theta}) > 0$. Moreover, for all $\theta \geq \underline{\theta}$, we have $\mathbf{r}(\theta) > 0$ 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 $\theta > \underline{\theta}$,

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

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 2(iii) that G^θ has strictly decreasing hazard rate, we have

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

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

Step 2: A symmetric monotone equilibrium of the original game.

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

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

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 effective bid satisfying $b_{ij} \geq \mathbf{b}(\underline{\theta})$, 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. □

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' making changes to rankings or rates in the TMS. Because demotions and auction events involve these kinds of changes, these timestamps provide candidates for the timing of demotion and auction events. We refer to the period between two consecutive timestamps of a shipper on a lane as a "date-range."

C.1 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

Table 3: 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.

multiple origin-destination pairs in close proximity into a single “lane”, using the same contract. Note that, in this latter case, the distance of the haul may vary from load to load. We use this to distinguish between two different pricing modes: linehaul pricing (where the payment is the same irrespective of distance) versus per-mile pricing. 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.

C.2 Primary status

From the fact that primary carriers are generally the first to receive shippers’ requests, we are able to identify primary carriers by observing which carriers tend to receive load offers first. Exceptions may arise, typically either due to pre-specified capacity constraints agreed upon by the shipper and the carrier or due to multiple primary carriers sharing primary status on the same lane. In such instances, we assign primary status to the carrier with the most first requests within a date-range.

C.3 Auction and demotion events

Our data do not include explicit 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.[58]

[58] There is a tradeoff in using this indicator. On the one hand, not using this indicator risks missing some

Figure 9: Number of identified auction and demotion events

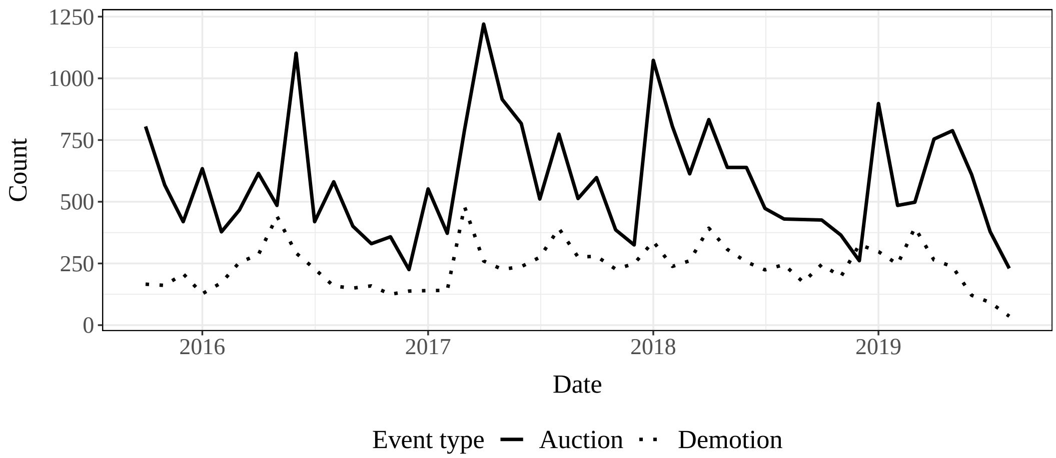

To identify demotion events, we identify instances in which the current primary carrier is replaced by another carrier within the same contract period (that is, between two auction events). Measurement errors in our constructed indicator of demotion events can come from measurement errors in our constructed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

Figure 9 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.

C.4 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 (NMFTA, 2021). We then map US DOT codes to carriers' registration at the Department of Transportation for the year 2020 (FMCSA, 2021). This method matches 90% of carriers in our microdata. We divide the matched carriers into

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.

Table 4: 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

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 E.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 is limited to asset-owner carriers. Table 4 provides summary statistics for this data set.

D Data representativeness

In our TMS data, we observe all of the interactions between 51 shippers and their contracted carriers. While this data contains 1.2 million transacted loads, this represents only about 0.1% of all US truckload contract transactions by value.[59]

To assess the representativeness of our sample, Figure 10 compares the distribution of fleet sizes of carriers in our data set with the distribution of fleet sizes of all registered carriers. While this comparison shows that the carriers appearing in the CHR data set tend to be large relative to the population of all carriers, the CHR data set is not necessarily unrepresentative of the population of carriers in long-term relationships. Large carriers are more likely to transact through long-term relationships, while small carriers are more likely to transact in the spot market.[60] Moreover, Figure 11 shows that the rates of rejection by primary carriers over time are quite similar across carriers of different sizes. This analysis of heterogeneity within our data set suggests that any overrepresentation of large carriers in our sample is unlikely to bias our results since—at least in the aggregate—carrier behavior within relationships is relatively uniform across the distribution of carrier sizes.

[59]This fraction is calculated based on the estimated annual trucking contract market in the U.S. at $288 billion (Holm, 2020*a*).

[60]See, for example, Holm (2020*a*): "Larger carriers tend to have more contract exposure, while smaller fleets generally have less...."

Figure 10: Carrier size distributions: CHR data set versus all registered carriers

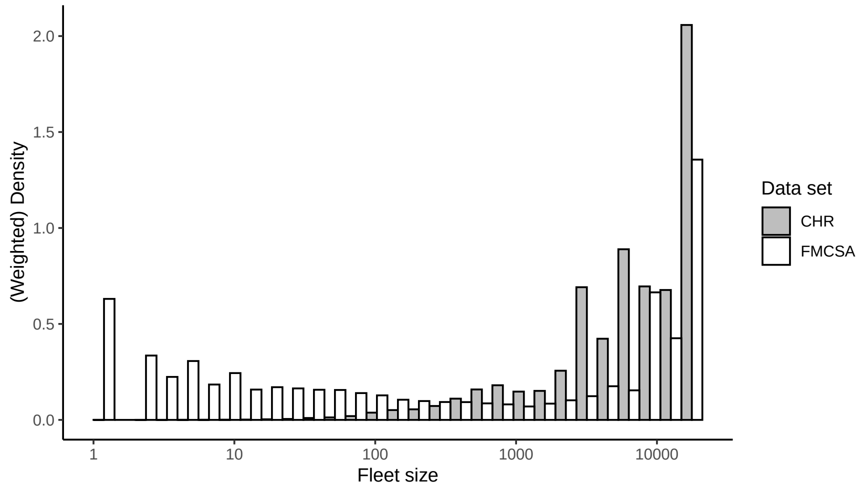

Note: This histogram compares the distribution of carriers' fleet size (number of trucks) in our CHR data set with that of the FMCSA Census of Carriers. Note that the x-axis has a log scale and that the y-axis shows the truck-weighted density. The latter means that rather than representing the fleet size distribution of *fleets*, these histograms represent the fleet size distribution of *trucks*. That is, if one were to draw a truck at random and observe the size of the fleet to which the truck belonged, the distribution of these fleet sizes would correspond to our histograms.

Figure 11: Monthly rejection rates by carrier size

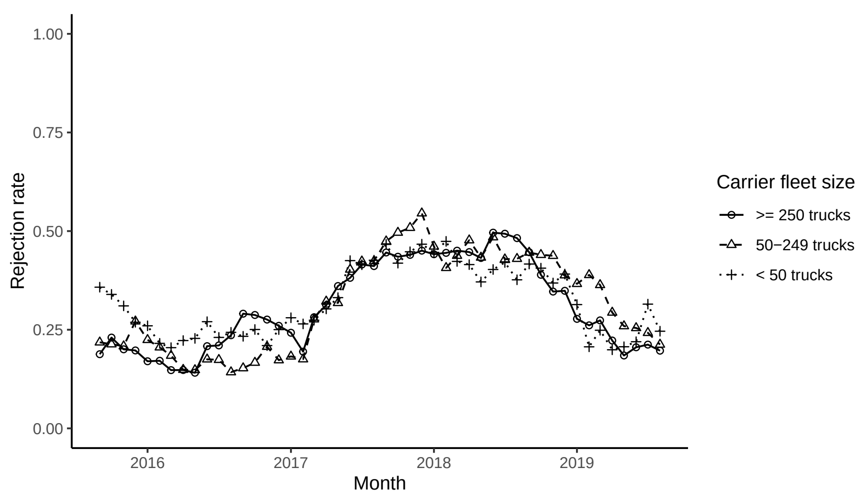

E Other empirical results

E.1 Sources of gains from long-term relationships

This section presents suggestive evidence on the match-specific gains from relationships (ψ, η) that play a central role in our welfare analysis. The primary purpose of presenting this evidence is to elucidate the sources of match-specific gains in the truckload setting. In addition, this evidence serves to support several of our modeling choices. In particular, it supports (i) our decision to treat match-specific gains as fixed throughout the relationship, (ii) our focus on acceptance—rather than other elements of carrier performance—as the relevant margin for moral hazard within relationships, and (iii) our decision to treat transactions with backup carriers as spot rather than relationship transactions.

To accomplish this, we compare carrier performance across different groups, analyzing differences between primary and backup carriers and between asset-owners and brokers. In addition to acceptance, which we emphasize as a key carrier decision margin throughout the paper, we exploit two additional measures of carrier performance: on-time delivery, which is a dimension of carrier reliability, and carrier detention time, which reflects the degree of shipper-carrier coordination.[61] To make these comparisons, we estimate regressions of the following form:

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

Here, $\text{Performance}_{ij\ell t}$ is one of our three performance measures (acceptance, on-time delivery, or detention time), and carrier's type is given by the interaction of asset ownership (asset-owner or broker) and relationship status (primary or backup). For each performance measure, we run both an (OLS) specification and a (FE) specification with shipper-carrier-lane fixed effects. Backup carriers are the excluded group in these (FE) specifications. We present the estimated coefficients α in Table 5 and discuss the interpretation of these coefficients below.

Sources of match-specific gains. We begin by comparing the OLS estimates of performance for primary asset-owners versus primary brokers. The idea behind this exercise is that, since brokers subcontract each load to an asset-based carrier through one-off transactions, transactions conducted via brokers serve as a useful proxy for spot transactions. We find that while brokers are about 10.1pp more likely to accept loads, asset-based carriers are

[61]Detention time refers to the duration that a driver is delayed at a shipper or receiver facility beyond the standard loading and unloading window. Such delays can result from inadequate dock space, delays in processing paperwork, or a lack of communication between the shipper and carrier. Detention time imposes significant costs on carriers, who are not typically compensated for detention time.

Table 5: Comparison of different groups of carriers across three performance measures

	Accept (0/1)		On-time (0/1)		Detention time (hours)	
	(OLS)	(FE)	(OLS)	(FE)	(OLS)	(FE)
<u>Asset owners</u>						
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)	
<u>Brokers</u>						
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): in addition to controls in the (OLS) specification, we also control for shipper-carrier-lane fixed effects. 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 detention time are only measured for accepted shipments.

10.8pp more likely to deliver loads on time and incur an average of 31 minutes less detention. The latter two observations support the existence of intrinsic gains in shipper-carrier relationships: shippers and asset-owner carriers that interact repeatedly have better service quality and shipper-carrier coordination.

Mechanisms underlying relationship gains. Better performance in relationships may arise through two channels: (i) a *selection* channel, where shippers and carriers anticipating repeated interactions invest more in finding high-quality matches, and (ii) an *incentive* channel, where the repeated nature of interactions itself fosters better performance. To disentangle these mechanisms, we compare the performance of primary versus backup carriers using the (OLS) specification—which captures both channels—and the (FE) specification—which isolates the incentive channel. For this exercise, we focus on asset owners.[62]

The OLS estimates show that primary carriers outperform their backup counterparts across all three performance measures: they have a 38pp higher acceptance rate, a 6.2pp higher likelihood of on-time delivery, and incur an average of 11 fewer minutes in deten-

[62]Primary and backup brokers exhibit very similar performance in both on-time delivery and detention time across the OLS and fixed effects (FE) specifications.

tion. However, when we include shipper-carrier-lane fixed effects, the differences in on-time delivery and detention time largely disappear. In contrast, about half of the difference in acceptance rates remains, with primary carriers still exhibiting a 14.3pp premium. This residual premium partly reflects carriers' responses to shippers' punishment schemes and partly reflects a coordination effect—primary carriers anticipate higher volume from contracted shippers and realign their networks accordingly to better accommodate these requests. We conclude that quality and coordination benefits—reflected in on-time delivery and detention time—are primarily driven by selection effects. Our model captures such benefits through match-specific gains ψ and η, which are held fixed throughout the relationship, with selection on these gains occurring at the auction stage.

Low acceptance as the moral hazard concern. For further evidence that acceptance is the key strategic decision margin for carriers, we examine how each performance measure responds to spot premiums. This responsiveness matters for two reasons: first, the spot premium shifts the carrier's outside option and thus the temptation to renege on relationships; and second, the spot premium is an equilibrium object shaped by market-level interactions between relationships and the spot market. We find that a one dollar increase in the spot rate relative to the contract rate results in an 18.5pp decrease in acceptance rate but no economically significant change in on-time delivery or detention time.

Backup similar to spot. Finally, we compare the performance of backup asset-owners with that of backup brokers, based on the idea that the latter can serve as a proxy for spot transactions. We find that both groups exhibit similarly low acceptance rates—about half that of their primary counterparts. In terms of quality measures such as on-time delivery and detention time, backup asset-owners perform at an intermediate level, roughly halfway between primary asset-owners and brokers. This evidence underpins our treatment of backup transactions as analogous to spot transactions in the welfare analysis.

Another potential concern with this treatment for welfare analysis is that backup carriers might not face the same search costs for backup loads as spot carriers do for spot loads. However, this concern is alleviated by the fact that each backup carrier has a low probability of receiving and ultimately accepting a load offer; the fact that backup loads account for one-fifth of total routing-guide volume is largely due to the long tail of this routing guide.

E.2 Shipper's punishment scheme

Table 6 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,

Table 6: 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.

we simulate two sets of relationships, one with an initial rejection and one with an initial acceptance. We find that an initial rejection (as compared to an initial acceptance) reduces the expected number of requests the carrier would receive by 3%. This suggests that the shipper's punishment scheme is soft but nevertheless generates meaningful incentives.

Moreover, we find that when the current spot premium is high or when the shipper has large volume on a lane (i.e., 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 each rejection by the carrier affects the continuation probability of his relationship in many future periods. These findings on shippers' punishment scheme are consistent with findings in our companion paper Harris and Nguyen (2025).

E.3 Other welfare effects

In this section, we examine the heterogeneous welfare effects of different market institutions on various types of shippers and carriers. First, we analyze how the current institution—characterized by *fixed-rate* contracts and relational incentives—allocates relationship surplus between shippers and carriers in relationships of varying match quality. Second, we compare welfare changes for shippers and carriers, both within and outside of relationships, under two counterfactual scenarios: *index-priced* and *spot-only*, relative to the current *fixed-rate* scenario.

Figure 12: Volume-weighted expected surplus from long-term relationships

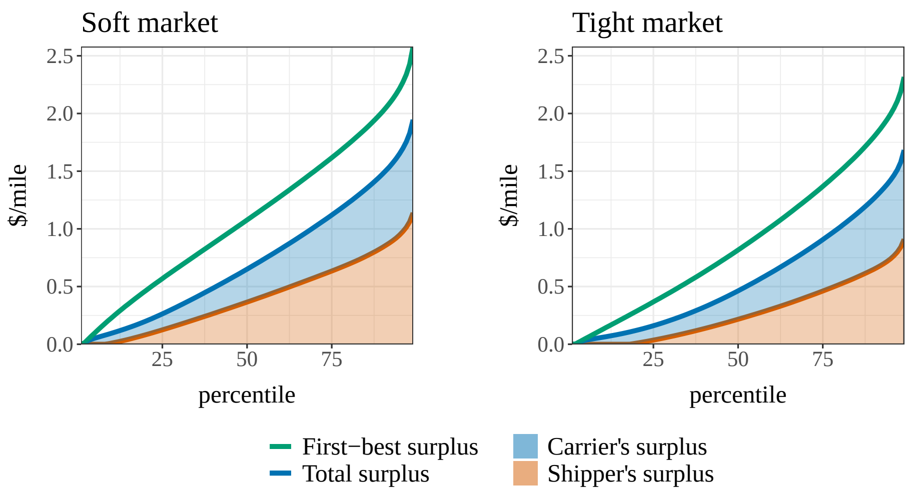

E.3.1 Relationship-level welfare effects: fixed-rate contracts

Figure 12 plots the volume-weighted expected surplus from current relationships (blue line) relative to spot transactions, showing how this surplus is split between shippers (red area) and carriers (blue area), alongside the relationship-level first-best surplus (green line). We find that the surplus of individual relationships is split roughly equally between the shipper and the carrier—a pattern that holds in both soft and tight markets.

The performance of current relationships relative to the relationship-level first-best benchmark is also broadly similar across soft and tight markets. However, there is substantial variation across relationships of different match quality. Specifically, a relationship at the 25th percentile of match quality captures only 45% of its potential relationship surplus, compared to 68% for a relationship at the 75th percentile. We draw two key conclusions from these findings. First, elevated rejection rates during tight markets are, at least in part, a reflection of higher (opportunity) costs faced by carriers, rather than purely indicative of relationship breakdowns. Second, under current contracts, the spot market disproportionately crowds out low-value relationships.

E.3.2 Market-level welfare effects: spot-only and index-priced

Next, we examine the welfare effects for different types of shippers and carriers under the *spot-only* and *index-priced* counterfactual scenarios. These effects operate through both price and non-price channels. Intuitively, a higher share of spot transactions reduces search costs, while improvements in allocative efficiency lead to lower spot prices. These changes in spot market conditions shape the outside options of shippers and carriers seeking to form relationships, thereby influencing how relationship surplus is divided between them.

Table 7: Effects on search costs and spot rates

	Fixed-rate	Spot-only	Index-priced
Soft market			
$\Delta\kappa$ ($/mile)		−0.75	0.09
Average spot rate ($/mile)	1.61	1.40	1.57
Tight market			
$\Delta\kappa$ ($/mile)		−0.43	0.22
Average spot rate ($/mile)	2.13	2.00	2.10

Table 7 presents the average spot rates under the *fixed-rate*, *index-priced*, and *spot-only* scenarios, along with changes in search costs in the latter two scenarios relative to the *fixed-rate* baseline. As expected, the *spot-only* scenario results in a substantial reduction in search costs—particularly in soft markets, where the baseline spot share is low. This is accompanied by a significant decline in spot rates. In contrast, the *index-priced* scenario increases search costs, but may still lead to lower spot rates, as the use of index-priced contracts contributes to improved allocative efficiency.

Figure 13 illustrates how these price and non-price effects translate into changes in the average per-period payoff ($/mile) for different groups of shippers and carriers under the *index-priced* and *spot-only* counterfactuals, relative to the baseline of *fixed-rate* contracts. We consider four groups: (i) shippers and (ii) carriers who would form relationships at baseline, and (iii) shippers and (iv) carriers who operate exclusively in the spot market.

First, we examine the *spot-only* counterfactual. 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 an increase in demand for spot loads and a reduction in search costs for spot loads.

Second, we examine the *index-priced* counterfactual—one in which we observed only a small change in aggregate welfare relative to the baseline (Figure 6). However, this aggregate finding masks substantial distributional consequences across agent types. We find that the *index-priced* counterfactual tends to favor shippers over carriers, with the largest welfare loss borne by carriers operating only in the spot market. These carriers are often small owner-operators, who might be forced to exit the market if faced with such loss.

A more subtle finding is that the widespread adoption of relationship-optimal *index-priced* contracts would not necessarily guarantee welfare gains for relationship participants either. 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 weak. The reason is that

Figure 13: Distributional consequences by different groups of shippers and carriers

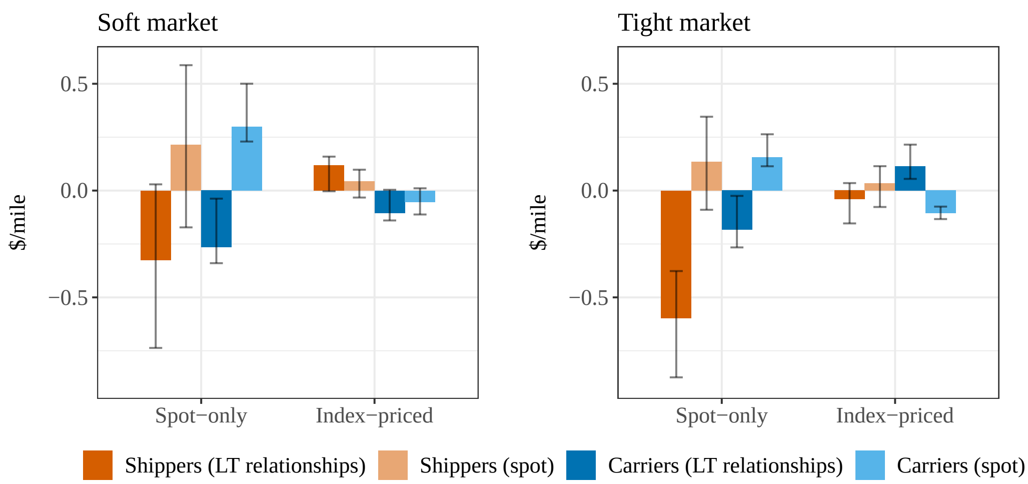

Note: Shippers' benefits from having their loads shipped are normalized to zero.

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 weakening their bargaining outcomes at the auction stage.

F Estimation details: Model primitives

F.1 Shippers' punishment schemes and instrumental objects

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.

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. We scale this daily discount rate by the frequency of shipments within each shipper-lane-auction to obtain the corresponding discount rate $\delta_{ia\ell}$.

Next, we describe how we estimate shippers' punishment scheme. 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 (6), 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 6 presents our estimates of shippers' strategy in both the linear (GMM) and Probit (IV Probit) specifications.

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 in estimating the distribution of the number of bidders 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.

F.2 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 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

$$\mathcal{L}\left((\mathbf{1}\{d_{ij\ell t} = \text{accept}\}, R_{ij\ell t-1}, \tilde{p}_{\ell t})_{t=1}^{T_{ij\ell}}; \sigma_\ell^\zeta, \tilde{\mu}_{ia\ell}^c, \sigma^c, \eta_{ij\ell} + p_{ij\ell} + \kappa_\ell\right) \tag{12}$$

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

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

$$\begin{aligned}&\Pr(d_{ij\ell t} = \text{accept}|R_{ij\ell t-1}, \tilde{p}_{\ell t}; \sigma_\ell^\zeta, \tilde{\mu}_{ia\ell}^c, \sigma^c, \eta_{ij\ell} + p_{ij\ell} + \kappa_\ell) \\ =&\Phi\left(\frac{\bar{p}(R_{ij\ell t-1}, \tilde{p}_t|\eta_{ij\ell} + p_{ij\ell} + \kappa_\ell; \sigma_\ell^\zeta, \tilde{\mu}_{ia\ell}^c, \sigma^c) - \tilde{p}_{\ell t}}{\sigma_\ell^\zeta}\right) \\ &\times \Phi\left(\frac{\bar{p}(R_{ij\ell t-1}, \tilde{p}_t|\eta_{ij\ell} + p_{ij\ell} + \kappa_\ell; \sigma_\ell^\zeta, \tilde{\mu}_{ia\ell}^c, \sigma^c) - \tilde{\mu}_{ia\ell}}{\sigma^c}\right).\end{aligned} \tag{13}$$

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 (3) 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.

Maximum-likelihood estimation of long-lasting relationships

In this first step, we use data only from 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.

Estimating cost parameters and extrapolating 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 (7). Second, we extrapolate these cost estimates to all relationships by running a flexible polynomial regression of these estimates on observed relationship and lane characteristics,

$$\widehat{\tilde{\mu}}^c_{ia\ell} = h(\tilde{\mathbf{X}}_{ia\ell}) + \tilde{\epsilon}^c_{ia\ell}. \tag{14}$$

Together, the extrapolated transformed costs from equation (14) and our estimates of cost parameters in equation (7) provide estimates $(\hat{\kappa}_\ell)_\ell$ of search costs 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 the previous, maximum-likelihood, step.

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 $\text{Rate}_\ell, \text{Volume}^{\text{spot}}_\ell, \text{Tight}_a$, 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 $\tilde{\mathbf{X}}_{ia\ell}$.

EM algorithm for pooling 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 using a conditional mixture model. Specifically, we assume that each shipper-lane-auction belongs to one of seven groups, $g = 1, ..., 7$, with the probability of being in each group g given by π_g. Conditioning on being in group g, the distribution of rent r of the winning carrier j in auction $ia\ell$ has a Normal distribution,

$$r_{ij\ell}|_{ia\ell\in g} \sim \text{Normal}(\beta^r_g \tilde{\mathbf{X}}_{ia\ell}, \sigma^r_g).$$

That is, we discipline the variation of carrier rents across auctions by allowing the conditional mean of each component distribution to vary linearly with observed relationship and lane characteristics; however, we do not allow these characteristics to influence the mixing probabilities $(\pi_g)_g$.

To estimate the mixing probabilities and the component distributions, we adapt an EM algorithm by Train (2008). First, before running the EM algorithm, we pre-calculate the likelihood contribution of each relationship at each grid point of carrier rent in $\{0.0, 0.1, ..., 5.0\}$. Since calculating the likelihood contribution is computationally intensive, performing this step only once and using linear interpolation from the above grid points in the integration step substantially speeds up our EM algorithm.

Except for this interpolation procedure, the EM algorithm is standard. In the M-step, we integrate the likelihood contribution in equation (12) over the mixture of carrier rents implied by the estimated parameters from the previous E-step. With the likelihood at each grid point already calculated, we only need to update the intepolation weights on these points to complete the integration. In the E-step, we update the mixing probabilities $\{\pi_g\}_g$, the mean-shifting parameters $\{\beta_g^r\}_g$, and variances $\{\sigma_g\}_g$ of the component distributions. We iterate between these two steps until convergence.

The output of this procedure 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 $\underline{r}_{ia\ell}$.

F.3 Estimating 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 above (see F.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.

F.4 Estimating 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 F.2. Given these conditional distributions and the bidding function from F.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).

G Estimation details: Welfare analysis

Since our relationship data covers only a subset of shippers, we cannot reliably estimate aggregate relationship demand at the KMA-to-KMA level. To address this limitation, we aggregate lanes into seven clusters using K-means clustering based on the following lane characteristics: the average spot rate (Rate_ℓ), the log of average spot volume (Volume_ℓ), and the average distance (Distance_ℓ). The selection of these characteristics—which are equilibrium objects shaped by underlying demand and supply factors—reflects our intention to aggregate lanes in a way that captures these underlying demand and supply factors (Bonhomme and Manresa, 2015). Figure 14 illustrates these clusters. The left panel plots the log of the relationships' average spot volume (horizontal axis) and the relationships' average spot rate (vertical axis), with the colors indicating the cluster to which each relationship belongs. The right panel plots the median cost (vertical axis) and the log of predicted trade flows (horizontal axis) for each cluster. The size of each cluster's point represents the number of relationships within that cluster. These scatter plots suggest that the equilibrium spot price and volume effectively separate the underlying demand and supply factors. Finally, we split each cluster in two according to market phase with the year 2016 representing a soft market and 2018 a tight market. We perform welfare analysis separately for each of these 14 resulting clusters of lanes and market phase (year).

Figure 14: Seven clusters of lanes identified by K-means clustering

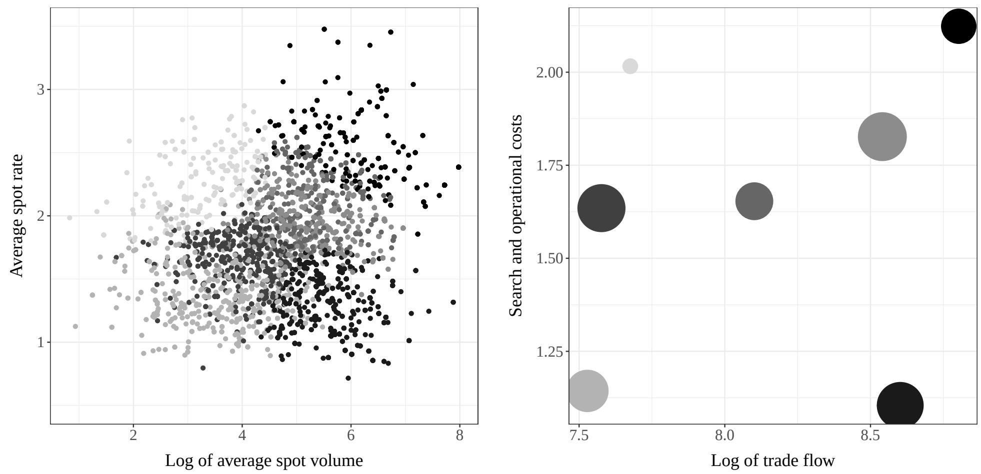

In performing our welfare analysis, we make the following assumptions: First, in period $t = 0$, there are L^* shipper-carrier relationships formed and C^* carriers in the spot market. Second, as relationships dissolve over time, shippers and carriers in those relationships join the spot market. This process, dictated by the shippers' punishment, decreases the number of relationships to L_t and increases the spot capacity to C_t in period t, while keeping fixed total carrier capacity, i.e., $L_t + C_t = L^* + C^*$.[63] Finally, exogenous shock D_t^* contributes to direct spot demand D_t. Our welfare analysis involves two steps: First, we estimate the underlying aggregate demand and capacity $(L^*, C^*, D_t^*)_{t=1}^{52}$. Second, we calculate welfare in the current and alternative institutions while keeping these underlying factors fixed.

G.1 Estimation of aggregate demand and capacity

For each cluster of lanes and each market phase, we recover the tuple $(L^*, C^*, D_t^*)_{t=1}^{52}$ of initial relationships, capacity, and shocks to spot demand by leveraging (i) observed prices and volumes in both relationships and the spot market and (ii) our estimates of match-specific gains from relationships and of carriers' costs. We recover $(L^*, C^*, D_t^*)_{t=1}^{52}$ by methods of moments—specifically, by matching the model-predicted paths of relationship demand, relationship volume, and spot volume $(L_t, L_t^{\text{primary}}, S_t)_{t=1}^{52}$ to their empirical counterparts.

To derive the model-predicted moments, we proceed in two steps. First, each guess (L^*, C^*) of initial relationships and capacity, together with our model estimates in Section 6 and observed prices, allows us to forward-simulate a path $(L_t, L_t^{\text{primary}}, S_t)_{t=1}^{52}$ of relationship demand, relationship volume, and spot volume. Recall that $\mathcal{H}_{t-1} = (R_{ij,t-1}, \eta_{ij} + p_{ij}, \psi_{ij} -$

[63]In reality, new relationships are formed over time and carriers can adjust their capacities by buying and selling trucks. To help with the estimation of market dynamics, we make the simplifying assumptions that all relationships are created in period 0 and capacity is fixed within a market phase.

$p_{ij})_{ij}$ captures the relevant information about relationship stock at the end of period $t-1$. At the beginning of period t, shippers terminate some of these relationships—a process influenced by carriers' rejection indices and the current spot rate—leaving

$$L_t(\tilde{p}_t|\mathcal{H}_{t-1}) = \underbrace{\int_{ij\in\mathcal{H}_{t-1}} (1-\sigma_0(R_{ij,t-1},\tilde{p}_t))}_{\text{Surviving relationships}} \tag{15}$$

relationships.[64] Carriers in terminated relationships join the spot market, increasing spot capacity to

$$C_t(\tilde{p}_t|\mathcal{H}_{t-1}) = C_{t-1} + \underbrace{\int_{ij\in\mathcal{H}_{t-1}} \sigma_0(R_{ij,t-1},\tilde{p}_t)}_{\text{Newly terminated relationships}} . \tag{16}$$

Given the relationships that persist into period t, the relationship and spot volumes realized at the end of this period further depend on carriers' decisions. Recall that a primary carrier's decision to accept a load within his relationship, serve the spot market, or remain idle depends on the relationship stock only through his full compensation. Consequently, in aggregate, the decisions of carriers within relationships can be summarized by the probability distribution of full compensation across carriers, denoted by $\mu(\bar{p}|\tilde{p}_t,\mathcal{H}_{t-1})$. The relationship volume and spot volume realized in period t are

$$L_t^{\text{primary}}(\tilde{p}_t|\mathcal{H}_{t-1}) = L_t(\tilde{p}_t|\mathcal{H}_{t-1}) \int_{\tilde{p}_t}^{\infty} F(\bar{p}-\kappa_\ell) d\mu(\bar{p}|\tilde{p}_t,\mathcal{H}_{t-1}) \tag{17}$$

and

$$S_t(\tilde{p}_t|\mathcal{H}_{t-1}) = \left(L_t(\tilde{p}_t|\mathcal{H}_{t-1})\mu(\tilde{p}_t|\tilde{p}_t,\mathcal{H}_{t-1}) + C_t(\tilde{p}_t|\mathcal{H}_{t-1})\right) F(\tilde{p}_t-\kappa_\ell) \tag{18}$$

respectively. That is, relationship carriers whose full compensation exceeds both the current spot rate and the sum of operational and search costs serve relationship loads. Carriers whose full compensation is lower than the current spot rate serve spot loads if and only if their costs are sufficiently low. Matching the forward-simulated path $(L_t, L_t^{\text{primary}}, S_t)_{t=1}^{52}$ from equations (15), (17), and (18) to their empirical counterparts pins down the initial relationships and spot capacity (L^*, C^*).

Second, we pin down the exogenous demand shock D_t^* in each period from the market

[64] Here, with some abuse of notation, we write $ij \in \mathcal{H}_{t-1}$ for a relationship that is maintained at the end of period $t-1$. Similarly, we write $ij \in L_t(\tilde{p}_t|\mathcal{H}_{t-1})$ for a relationship that survives the shipper's punishment strategy at the beginning of period t.

clearing condition. Specifically, this demand shock, together with all terminated relationships, contributes to direct spot demand,

$$D_t(\tilde{p}_t|\mathcal{H}_{t-1}) = \underbrace{L^* - L_t(\tilde{p}_t|\mathcal{H}_{t-1})}_{\text{All terminated relationships}} + D_t^*; \tag{19}$$

in equilibrium, the total demand from relationships and spot equals the total realized volume,

$$L_t(\tilde{p}_t|\mathcal{H}_{t-1}) + D_t(\tilde{p}_t|\mathcal{H}_{t-1}) = L_t^{\text{primary}}(\tilde{p}_t|\mathcal{H}_{t-1}) + S_t(\tilde{p}_t|\mathcal{H}_{t-1}). \tag{20}$$

That is, exogenous demand shocks help rationalize the residual variation in spot rates that is not explained by relationship dynamics.

Finally, we explain how to derive the empirical counterparts of our volume moments, $(L_t, L_t^{\text{primary}}, S_t)_{t=1}^{52}$, from our merged relationship and spot data. For each period t, we obtain the total number of loads $\hat{L}_t$ requested by the shippers in our sample, the loads $\hat{L}_t^{\text{primary}}$ accepted by the primary carriers, the loads $\hat{L}_t^{\text{backup}}$ accepted by the backup carriers, as well as the residual loads that must be fulfilled in the spot market. From our spot data, we obtain the number of spot loads $\hat{S}_t'$, scaled 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. This decision to pool backup and spot loads is based on the evidence in Appendix E.1, which suggests that the performance of backup carriers is more similar to that of spot carriers than to that of primary carriers. Thus, our empirical moment for realized spot volume is $\hat{S}_t = \hat{S}_t' + \hat{L}_t^{\text{backup}}$. When we pool loads fulfilled by backup carriers with spot volume, the spot share under the current institution becomes 38%.

G.2 Welfare of the current institution

As explained above, a tuple $(L^*, C^*, D_t^*)_{t=1}^{52}$, together with the observed path $(\tilde{p}_t)_{t=1}^{52}$ of spot rates, contract rates, and model estimates of relationship and cost parameters, allows us to forward-simulate the dynamics of relationships and their interactions with spot movements. We define welfare in period t in the current institution as

$$\begin{aligned} W_t^0 = &\int_{ij \in L_t(\tilde{p}_t|\mathcal{H}_{t-1})} \mathbf{E}_F\left[(\theta_{ij} - c_{jt})\mathbf{1}\{\bar{p}(R_{ij,t-1}, \tilde{p}_t|\eta_{ij} + p_{ij}) \geq \max\{\tilde{p}_t, c_{jt} + \kappa\}\}\right] \\ &+ \left[\underbrace{C_t(\tilde{p}_t|\mathcal{H}_{t-1})}_{\text{direct spot capacity}} + \underbrace{L_t(\tilde{p}_t|\mathcal{H}_{t-1})\mu(\bar{p}|\tilde{p}_t, \mathcal{H}_{t-1})}_{\text{overflow capacity}}\right](-\kappa - \mathbf{E}_F[c_{jt}|c_{jt} \leq \tilde{p}_t - \kappa])F(\tilde{p}_t - \kappa). \end{aligned}$$

The first component represents welfare gains from loads accepted within relationships. The second component represents welfare gains from spot transactions, which are fulfilled both by direct spot carriers and carriers whose full compensations from accepting relationship loads are less than the current spot rate. Crucially, our welfare measure normalizes the gains to shippers from having their loads fulfilled by spot carriers to zero. Since total shipment demand remains fixed across our counterfactual analyses, this normalization does not affect our welfare comparison.[65]

G.3 Welfare of counterfactual institutions

We compare welfare under the current institution with that achieved under counterfactual institutions that vary the split between relationship and spot transactions while keeping the underlying demand and capacity fixed. To focus on the relationship-spot split, we derive the constrained first-best welfare that could be achieved for any given share of spot volume in total volume. We will show that this is a linear programming problem. Moreover, we show that this constrained first-best welfare can be achieved by pairing index-priced contracts with an appropriate tax on relationship transactions.

To build intuition, consider the sub-problem of allocating carriers to relationships and spot in each period while keeping the path $S^t = (S_t)_{t=1}^{52}$ of spot volumes fixed. Fix the set of relationships L^* and spot carriers C^*. With some abuse of notation, we write $j \in L^*$ for a carrier with a relationship, whose match quality with shipper $i(j)$ is $\theta_{i(j)j}$, and we write $j \in C^*$ for a spot carrier. 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 $j \in C^*$. The constrained first-best allocation in period t, for a given path of spot volumes, solves

$$\begin{aligned}\overline{W}_t(S^t) = \max_A & \int_{j \in L^*} \mathbf{1}\{A(j) = \text{accept}\}(\theta_{i(j)j} - c_{jt}) \\ & + \int_{j \in L^* \cup C^*} \mathbf{1}\{A(j) = \text{spot}\}\left(-c_{jt} - \kappa\left(\bar{S}\right)\right) \\ \text{s.t.} & \int_{j \in L^*} \mathbf{1}\{A(j) = \text{accept}\} = L^* + D_t^* - S_t \\ & \int_{j \in L^* \cup C^*} \mathbf{1}\{A(j) = \text{spot}\} = S_t,\end{aligned}$$

[65] Note that by adjusting the split between relationship and spot transactions in counterfactual analyses, we allow spot rates to vary across different institutions. In principle, changing spot rates could influence the aggregate demand for trucking services. We shut down this channel to focus on allocative efficiency between relationship and spot transactions. As we later show, allocative efficiency could be improved by levying a tax on relationship transactions. Similarly, concerns about the extensive margin of trucking services could be addressed through a tax or subsidy on the industry as a whole.

where $\bar{S} = \frac{1}{52}\sum_{t'=1}^{52} S_{t'}$.

Let λ_t^1 and λ_t^2 be the Lagrange multipliers associated with the constraints on relationship and spot volumes, respectively, in period t. The Lagrangian of the above constrained optimization problem is

$$\int_{j\in L^*} \left[\mathbf{1}\{A(j) = \text{accept}\}(\theta_{i(j)j} + \lambda_t^1 - c_{jt}) + \mathbf{1}\{A(j) = \text{spot}\}(\lambda_t^2 - c_{jt} - \kappa(\bar{S}))\right] \\ + \int_{j\in C^*} \mathbf{1}\{A(j) = \text{spot}\}(\lambda_t^2 - c_{jt} - \kappa(\bar{S})).$$

The solution to this problem involves each $j \in L^*$ choosing the maximal in $\{\theta_{i(j)j} + \lambda_t^1 - c_{jt}, \lambda_t^2 - c_{jt} - \kappa(\bar{S}), 0\}$ and each $j \in C$ choosing the maximal in $\{\lambda_t^2 - c_{jt} - \kappa(\bar{S}), 0\}$. That is, the constrained first-best allocation is achieved by combining index-priced contracts with a period-specific corrective tax on relationship transactions. In equilibrium, the spot rate equals λ_t^2 and the corrective tax equals $\lambda_t^2 - \lambda_t^1$.

Given the above sub-problem, the constrained first-best welfare associated with spot share $s \in [0,1]$ can be written as

$$\begin{aligned} \overline{W}(s) = &\max_{S^t} \sum_{t=1}^{52} \overline{W}_t(S^t) \\ &\text{s.t. } \frac{\frac{1}{52}\sum_{t'=1}^{52} S_{t'}}{L^* + \frac{1}{52}\sum_{t'=1}^{52} D_{t'}^*} = s \\ &\quad 0 \le S_t \le L^* + D_t^*, \forall t. \end{aligned}$$

Thus, the market-level first best equals $\max_{s\in[0,1]} \overline{W}(s)$.

Besides this first-best benchmark, we will zoom in onto two special cases: (i) the *spot-only* scenario and (ii) the *index-priced* scenario with zero tax on relationship transactions. First, in the *spot-only* scenario, each carrier either serves the spot market or remains idle, resulting in aggregate welfare of

$$\begin{aligned} W^1 = &\frac{1}{52}\sum_{t=1}^{52}(L^* + D_t^*)(-\kappa^1 - \mathbf{E}_F[c_t | c_t \le \tilde{p}_t^1 - \kappa^1]) \\ &\text{s.t. } \kappa^1 = \kappa\left(\frac{1}{52}\sum_{t=1}^{52}(L^* + D_t^*)\right) \\ &\quad L^* + D_t^* = (L^* + C^*)F(\tilde{p}_t^1 - \kappa^1), \forall t. \end{aligned}$$

Here, κ^1 is the equilibrium search cost and $(\tilde{p}_t^1)_{t=1}^{52}$ is the market-clearing path of spot rates.

Second, consider the *index-priced* scenario with zero tax. Denote by κ^2 the equilibrium search cost in this scenario and by $(\tilde{p}_t^2)_{t=1}^{52}$ the equilibrium price path. In this case, only relationships with $\theta \geq -\kappa^2$ are formed, and carriers in these relationships never go to the spot market; they accept relationship loads if and only if $\theta + \tilde{p}_t^2 \geq c_{jt}$. Carriers in potential relationships with $\theta < -\kappa^2$ join the spot market, giving rise to new spot capacity C_2^*; these carriers serve the spot market as long as $\tilde{p}_t^2 \geq c_{jt} + \kappa^2$. Aggregate welfare in this case is

$$
\begin{aligned}
W^2 = \frac{1}{52}\sum_{t=1}^{52} & \left(\begin{array}{c} L^* \int_{\theta \geq -\kappa^2} F(\theta + \tilde{p}_t^2)(\theta - \mathbf{E}_F[c_t | c_t \leq \theta + \tilde{p}_t^2]) d[G^\theta]^{1:N}(\theta) \\ +C_2^* F(\tilde{p}_t^2 - \kappa^2)(-\kappa^2 - \mathbf{E}_F[c_t | c_t \leq \tilde{p}_t^2 - \kappa^2]) \end{array} \right) \\
\text{s.t. } \kappa^2 &= \kappa \left(\frac{1}{52} \sum_{t=1}^{52} C_2^* F(\tilde{p}_t^2 - \kappa_2) \right) \\
C_2^* &= C^* - L^*[G^\theta]^{1:N}(-\kappa) + L^*[G^\theta]^{1:N}(-\kappa^2) \\
L^* + D_t^* &= L^* \int_{\theta \geq -\kappa^2} F(\theta + \tilde{p}_t^2) d[G^\theta]^{1:N}(\theta) + C_2^* F(\tilde{p}_t^2 - \kappa^2), \forall t.
\end{aligned}
$$

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 the full relationship surplus, denoted by $\text{Surplus}(\theta_{ij}|\tilde{p}_0)$, is realized, and that 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

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