

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

Adam Harris[†] Thi Mai Anh Nguyen[‡]

Job Market Paper

January 30, 2023

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

Abstract: Long-term informal relationships play an important role in the economy, capitalizing on match-specific efficiency gains and mitigating incentive problems. However, the prevalence of long-term relationships can also lead to thinner, less efficient spot markets. We develop an empirical framework to quantify the market-level tradeoff between long-term relationships and the spot market. We apply this framework to an economically important setting—the US truckload freight industry—exploiting detailed transaction-level data for estimation. At the relationship level, we find that long-term relationships have large intrinsic benefits over spot transactions. At the market level, we find a strong link between the thickness and the efficiency of the spot market. Overall, the current institution performs fairly well against our first-best benchmarks, achieving 44% of the relationship-level first-best surplus and even more of the market-level first-best surplus. The findings motivate two counterfactuals: (i) a centralized spot market for optimal spot market efficiency and (ii) index pricing for optimal gains from individual long-term relationships. The former results in substantial welfare loss, and the latter leads to welfare gains during periods of high demand.

This research was made possible by the many people in the truckload freight industry who generously shared their time, insights, and data with us. We are also very grateful to Angi Acocella, Nikhil Agarwal, Chris Caplice, Glenn Ellison, Bob Gibbons, Stephen Morris, Nancy Rose, Tobias Salz, and Mike Whinston for their advice and comments. We also thank all participants in the MIT Industrial Organization Lunch and Seminar, the MIT Organizational Economics Lunch and Seminar, and the MIT Freight Lab for their comments and feedback. We acknowledge the support of the George and Obie Shultz Fund. While this paper makes use of data provided by CH Robinson and DAT, the opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of CH Robinson or DAT.

Harris: This material is also based upon work supported by the National Science Foundation Graduate Research Fellowship under Grant No. 1745302. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

Nguyen: This material is based upon work supported by the Jerry A. Hausman Graduate Dissertation Fellowship.

[†]MIT Department of Economics. Email address: asharris@mit.edu.

[‡]Corresponding author. MIT Department of Economics. Email address: anhng@mit.edu.

1 Introduction

Long-term relationships are a ubiquitous feature of the economy (Macaulay, 1963; Macchiavello, 2022), and in many settings, these relationships coexist and interact with other forms of transactions (Allen & Wittwer, 2021; Macchiavello & Morjaria, 2015). While the ubiquity of relationships suggests that they benefit the participating parties, it does not rule out the possibility that relationships exert negative externalities on the rest of the market. Long-term relationships are individually rational, but are they socially efficient?

We develop an empirical framework to quantify the market-level value of long-term relationships and alternative institutions in the setting of the US for-hire truckload freight industry, one in which long-term relationships coexist with spot marketplaces. Combining elements from the auction and dynamic discrete choice literature, our model captures both the formation of and interactions within long-term relationships. The demand side (shippers) and the supply side (carriers) form relationships via auctions, facilitating the formation of relationships with high match quality. Within relationships, carriers face a temptation to defect to the spot market when spot rates are high, and shippers use a relational incentive scheme to mitigate this problem. At the market level, we allow for two-way crowding-out effects between long-term relationships and the spot market (Kranton, 1996), finding strong effects in both directions. On the one hand, the spot market creates a moral hazard problem in long-term relationships, crowding out low-value relationships. On the other hand, the formation of long-term relationships results, in aggregate, in a thin and substantially less efficient spot market. We quantify this tension in two counterfactual exercises: (i) centralizing all transactions into a spot market for maximal spot market thickness and (ii) using index-priced contracts in long-term relationships to resolve the moral hazard problem.

The US for-hire truckload freight industry is an important economic setting and one in which long-term relationships play a central role. In 2019, this industry generated revenues equivalent to 0.8% of US GDP and transported 72% of domestic shipments by value, playing an integral part in US domestic trade. Consequently, the outcomes in this industry have implications for supply chains and the goods economy as a whole. In this market, 80% of total transacted volume is accounted for by long-term relationships. Relative to spot transactions, long-term relationships may allow their participants to enjoy a variety of possible benefits: more reliable service; more seamless loading, docking, insurance, and payment processing; better planning; or lower transaction costs.¹ While our empirical framework does not decom-

¹In this setting, shippers (the demand side) may care about reliability of service, efficiency at loading and docking, or ease of communication; carriers (the supply side) may care about the ability to find backhauls, familiarity with facilities, or promptness of payment. See Hubbard (2001) and Masten (2009) for a discussion on non-price factors that matter in this industry.

pose relationship benefits into specific channels, we quantify the total benefits of relationships, the extent to which these benefits are realized, and how the realized benefits are shaped by market structure.

The US for-hire truckload freight industry offers an ideal empirical setting for studying long-term relationships and their interactions with the spot market. In this setting, shippers are firms that demand transportation service on some origin-destination pair (“lane”). Carriers are transportation firms that are hired to provide such service. Within a long-term relationship, a contract fixes the price (“rate”) but not the volume. This leaves room for relational incentives to govern transactions. In many settings, a lack of data on interactions within informal relationships poses a significant obstacle to studying such relationships. Our setting, however, is one in which such interactions between shippers and carriers leave a digital record. Specifically, shippers use transportation management systems (TMS) to automate many aspects of their relationships with carriers; such systems record shippers' offers, carriers' responses, and the status of relationships at each point in time. Our study uses an anonymized panel of relationships of shippers that use one of the TMSs. To capture the outside option of long-term relationships, we obtain data on spot rates and volumes at fine spatial and temporal granularity. The combined data set allows us to have both a microscopic view of individual relationships and a bird's-eye view of their aggregate effects on the spot market.

We start in Section 4 with key patterns in the data. First, we find large heterogeneity in shippers' and carriers' behaviors across relationships after conditioning on contract and spot rates. This suggests the role of non-price factors in generating match-specific gains 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.² A potential explanation for this result is that lanes with higher total market demand have the potential for achieving higher spot market thickness, which, in turn, increases the relative attractiveness of spot arrangements. The latter link means that long-term relationships crowd out the spot market by making it thinner and less attractive.

In Section 5, we develop a model that allows us to quantify the tension between realizing relationships' match-specific gains and maintaining spot market efficiency. The modeling of individual relationships combines elements of models from the auction and dynamic discrete choice literature. In our model, each relationship consists of two stages. In the first stage, the shipper holds an auction to select a carrier with whom to form a relationship. In the second stage, the shipper and the winning carrier interact in a repeated game. In each period of this

²While Hubbard (2001) exploits equilibrium supply-side variation, we exploit the predicted trade flows between different states of the US (Caliendo, Parro, Rossi-Hansberg, & Sarte, 2018) as an exogenous source of demand-side variation.

game, the shipper decides whether to terminate the relationship or maintain the relationship and offer a load; the carrier decides whether to accept the load, reject it for a spot offer, or reject it to remain idle. Motivated by the data patterns established in Section 4, we allow for two-sided match-specific gains from relationships and capture the link between spot market thickness and efficiency via a search cost that the carrier incurs from servicing the spot market. Motivated by evidence from Harris and Nguyen (2021), we model the shipper as using an incentive scheme that conditions the probabilistic termination of relationships on carriers' past rejections. The carrier responds optimally in each period, taking into account the current and future compensation for an accepted offer, the compensation and search cost in the spot market, and the operational cost for the current period. At the market level, long-term relationships and the spot market interact in two ways. First, spot rates are determined in equilibrium, absorbing both direct spot demand and rejected offers from long term relationships. Second, search costs on the spot market are determined endogenously by the equilibrium spot volume.

In Section 6, we show that shippers' and carriers' primitives are nonparametrically identified from their behaviors in the auction and repeated game. Our identification strategy illustrates how insights into both the formation of and interactions within relationships help recover a rich set of model primitives. In three sequential steps, we identify carriers' primitives from their dynamic play in the repeated game. The first step identifies the distribution of the sum of search and operational costs. The argument for the identification of this distribution is motivated by the following thought experiment: If a carrier could only decide between “spot” and “idle”, the probability that this carrier chose “spot” over “idle” at different spot rates would trace out the distribution of the sum of search and operational costs. However, the empirical challenge in our setting is that we only observe whether the carriers in long-term relationships—who decide between “accept”, “spot”, and “idle”—accept or reject. To overcome this challenge, we develop a support-based argument that relates the unobserved decision margin (between “spot” and “idle”) to the observed decision margin (between “accept” and “reject”), thereby locally recreating the hypothetical scenario. The second step decomposes search and operational costs from their sum. Here we establish causality between search costs and spot market thickness by exploiting the predicted trade flows between different states of the US (Caliendo, Parro, Rossi-Hansberg, & Sarte, 2018) as a demand shifter. The third step identifies carriers' match-specific gains from their acceptance probability, observed prices, and the cost parameters identified in the previous steps. Intuitively, a carrier's match-specific gain can be inferred from the level of spot rate at which this carrier's acceptance becomes responsive to spot rates.

Next, we exploit the fact that relationships are formed via auctions to identify the distribution of shippers' match-specific gains. Intuitively, the way that carriers' match-specific gains are reflected in their bidding is determined by how contract rates split total match-specific

gains into carriers' rents and shippers' rents. Building on this observation, we derive, under empirically plausible conditions, equilibrium conditions that give rise to a monotone mapping between carriers' rents (whose components are either observed or already identified) and shippers' rents (which we need to identify). In the spirit of Guerre, Perrigne, and Vuong (2000), we pin down this monotone mapping from the first-order condition of carriers' bidding in the space of carriers' rents. Finally, identified rents and observed contract rates recover match-specific gains.

In Section 7, we present our estimates of the model primitives, showing that the key tension in our setting is between the large benefits of long-term relationships to the participating parties and their substantial negative externalities on the spot market. On the one hand, we find that each shipper and carrier in long-term relationships enjoys an average premium over a spot transaction of 58% and 10%, respectively, for each realized transaction. Moreover, current fixed-rate contracts and relational incentive schemes capture these potential premiums fairly well. Specifically, the current relationships achieve, on average, 44% of the relationship-level first-best surplus, with large heterogeneity across relationships with different match quality. Underlying this heterogeneity is the fact that, in relationships with lower match quality, carriers' moral hazard is more severe, and relational incentives are less effective at mitigating this problem. In other words, the spot market crowds out in particular long-term relationships with low match quality. On the other hand, relationships' formation and high performance result in a thin spot market with substantially higher search costs. We estimate that doubling the thickness of the spot market on a lane reduces search costs by an amount equivalent to reducing operational costs by 29%.

Motivated by these findings, Section 8 evaluates two counterfactual institutions: (i) a centralized spot market for optimal spot market efficiency and (ii) individually first-best contracts for optimal performance in long-term relationships. Comparing the current institution to the first counterfactual institution suggests that the current dominance of long-term relationships is not due to a coordination failure to form a thick spot market but instead due to the considerable benefits of long-term relationships. Specifically, we find that centralizing all transactions into a spot market results in substantial welfare loss, equivalent to about 25% of the median operational cost. In theory, a centralized spot market has two potential sources of gains: (i) reduction in search costs and (ii) improvements in allocative cost efficiency. However, our estimates suggest that both of these gains are small. One reason is that while substantially reducing search costs on the spot market, a centralized spot market increases search costs for those who would otherwise be in relationships and not incur any search costs. Overall, cost reductions from centralizing all transactions are not nearly enough to compensate for the complete loss of match-specific gains from long-term relationships.

In the second counterfactual exercise, we replace fixed-rate contracts with index-priced contracts designed to achieve the first-best welfare for individual relationships. Comparing the market-level performance of these contracts highlights the key tradeoff in our setting: any attempt to improve the performance of long-term relationships would worsen their negative externalities on the spot market. Specifically, while these contracts increase the realized relationship benefits by 11% to 28%, such gains are roughly offset by a substantial increase in search costs in the spot market and a reduction in allocative cost efficiency. Only in periods of high demand, when fixed-rate contracts face serious moral hazard, does the former effect dominate, leading to welfare gains from index-priced contracts. Overall, both fixed-rate and index-priced contracts perform fairly well at the market level, achieving at least 40% of the market-level first-best surplus for medium trips of around five hundred miles and at least 60% of the market-level first-best surplus for long trips of around a thousand miles.

The paper proceeds as follows. Section 2 reviews the related literature. Section 3 presents institutional details. Section 4 describes our data and presents the data patterns that motivate our model. Section 5 describes our model. Section 6 explains our identification argument and the estimation procedure. Section 7 presents our estimates of key model primitives, and Section 8 presents counterfactual results.

2 Literature review

Our paper contributes to two empirical literatures—the literature on long-term informal relationships and the literature on trucking—with an empirical framework that combines tools from the auction and dynamic discrete choice literatures. We divide our literature review into three subsections relating to literatures on long-term relationships, trucking, and spot market efficiency.

Long-term informal relationships. The empirical literature on long-term informal relationships has developed a rich set of insights into the mechanisms through which relationships create value for participating parties and respond to external factors.³ We contribute to this literature by quantifying both the value of relationships to participating parties and the negative externalities of relationships on the spot market. Conceptually, our paper is most closely related to Kranton (1996), who uses a market equilibrium model to theorize two-way crowding-out

³For example, value creation in long-term informal relationships can arise from supply reliability (Adhvaryu, Bassi, Nyshadham, & Tamayo, 2020; Cajal-Grossi, Macchiavello, & Noguera, 2022), reputation building (Macchiavello & Morjaria, 2015), or relational adaptations (Barron, Gibbons, Gil, & Murphy, 2020). Relationships may terminate (Macchiavello & Morjaria, 2015) or restructure (Gil, Kim, & Zanarone, 2021) in the face of large shocks, and can be hampered by competition (Macchiavello & Morjaria, 2021).

effects between long-term relationships and the spot market. The first direction is, as argued by Baker, Gibbons, and Murphy (1994), that the spot market is the outside option of relationships, crowding out long-term informal relationships by making relational incentives harder to enforce. The other direction takes place at the market level. As relationships are formed, the spot market becomes thinner and less efficient; this, in turn, reinforces the relative attractiveness of relationships.⁴ To the best of our knowledge, we are the first to quantify this second direction.

Methodologically, we contribute an empirical framework that takes advantage of both the formation of and interactions within long-term relationships to recover a rich set of primitives. Specifically, we build on techniques in the dynamic discrete choice literature for both long panels (Rust, 1994) and short panels (Kasahara & Shimotsu, 2009) to recover model primitives on the carriers' side. Our problem is not standard in this literature but closely related to the literature on contracting with moral hazard (Perrigne & Vuong, 2011), in that payoff-relevant actions are not fully observed. Rather than relying on a mapping between unobserved actions and observables (Gayle & Miller, 2015) or the state transition process (Hu & Xin, 2021), we develop a support-based argument that relates the unobserved decision margin to the observed decision margin. We then adapt techniques from the empirical auction literature (Guerre, Perrigne, & Vuong, 2000) to recover primitives on the shipper' side, factoring in the equilibrium path of play in each potential relationship.

A few other papers quantify the value of long-term relationships, but they differ from our paper both methodologically and conceptually. Macchiavello and Morjaria (2015) exploit temporal variation in spot rates to bound the value of trading relationships in the Kenyan rose market. Our empirical strategy builds on their idea that variation in spot rates helps trace relationship value but differs in that we recover primitives that can be used for counterfactual analysis. A recent paper that recovers the primitives of relationships and performs counterfactual analysis is Bruges (2020), which studies the trading relationships in the manufacturing supply chain of Ecuador. This paper exploits the optimality conditions of dynamic contracting with exible monetary transfers to identify the distribution of buyer's type. Since relationships in our setting use fixed-rate contracts, we cannot apply similar methods. Instead, we exploit the formation of relationships via auctions and the rich dynamics within relationships to identify the distribution of two-sided match-specific gains from relationships. Startz (2021) is among the few papers, including ours, that take a market equilibrium approach and capture both

⁴Tunca and Zenios (2006) make a similar theoretical argument by examining the competition between procurement auctions and long-term relationships. The broad idea that market thickness can be self-filling is examined in other settings. For example, Ngai and Tenreyro (2014) show that thick-market effects can amplify the seasonality of the housing markets; in the setting of labor markets with match-specific quality, Elliott (2014) argues that thick-market effects give rise to multiplicity of search equilibria, suboptimal entry and market fragility.

the formation of and interactions within relationships. This paper quantifies the search and contracting frictions faced by Nigerian importers of consumer goods by exploiting importers' decisions to travel to the source country as a way of reducing both frictions. A key difference between our paper and Startz (2021) is that we identify and quantify the thickness-externality of long-term relationships on the spot market.

Trucking. Exploration of long-term contracts and spot arrangements in trucking dates back to Hubbard (2001), who finds that selection into spot transactions increases with market thickness, and Masten (2009), who argues that savings on transaction costs are an important driver for long-term contracts.⁵ Search costs in our model have a similar interpretation as transaction costs in Masten (2009); both increase the relative attractiveness of relationships. One additional insight from our paper is that these costs increase endogenously when more relationships are formed and the spot market becomes thinner. This insight also offers an additional explanation for the link between market thickness and the use of spot arrangements found in Hubbard (2001). On thicker lanes, there is more potential for spot market thickness, so search costs tend to be lower, which in turn increases the relative attractiveness of spot arrangements.

Since these papers, there have been significant improvements in how the trucking industry organizes shipper-carrier matching and how firms keep track of their interactions within relationships. These improvements have generated rich transaction-level data, from which our paper benefits. In our previous paper, Harris and Nguyen (2021), we establish an understanding of the nature and effects of dynamic incentives in long-term relationships in the US trucking industry. Our current work builds on such understanding, but the goal is to quantify the market-level tradeoff between long-term relationships and the spot market. We examine this tradeoff in two counterfactual exercises: (i) a centralized spot market for maximal spot market thickness and (ii) first-best contracts for optimal performance within long-term relationships. Our paper is the first to study two-way market-level interactions between long-term relationships and the spot market in a transportation setting. A series of papers in the transportation and logistics literature also uses the same data as our paper to study the effects of relationships on participating parties, examining reciprocity (Acocella, Caplice, & Sheff, 2020), factors that affect the value of relationships to carriers (Acocella, Caplice, & Sheff, 2022b), and potential Pareto improvements for participating parties from index pricing (Acocella, Caplice, & Sheff, 2022a). In the economics literature, Yang (2021) studies the home bias of truck drivers using

⁵A related strand of literature studies asset ownership in the trucking industry. For example, Baker and Hubbard (2003) study shippers' choices over private fleets versus for-hire carriers; Baker and Hubbard (2004) and Nickerson and Silverman (2003) study drivers' ownership of trucks. Another strand of the literature (Marcus, 1987; Rose, 1985, 1987; Ying, 1990) studies the effects of deregulation on the trucking industry. These papers provide an important historical context for the evolution of this industry.

a market equilibrium model but focuses exclusively on the spot market.

Spot market efficiency. The benefits and, to a lesser extent, the tradeoffs of centralizing all transactions into a spot platform have been examined in other settings. In other transportation markets, the benefits include the reduction of search costs from economies of density (Frechette, Lizzeri, & Salz, 2019) and the correction of spatial misallocation via platform pricing (Buchholz, 2022; Lagos, 2000, 2003). Tradeoffs include the platform exploiting its market power to extract surplus from both sides of the market (Brancaccio, Kalouptsi, Papageorgiou, & Rosaia, 2020; Rosaia, 2020). The key difference between these papers and ours is the central role of long-term relationships in our setting. Specifically, we model economies of density as the channel through which long-term relationships exert externalities on the spot market. We abstract from spatial misallocation, focusing instead on the (mis)allocation of transactions between long-term relationships and the spot market.

Multiple modes of transaction. The coexistence of multiple modes of transaction has been examined in various settings. Such coexistence could generate quality dispersion (Galenianos & Gavazza, 2017), and market frictions could shift market modality (Gavazza, 2010, 2011). In the setting of the liquefied natural gas industry, Zahur (2022) shows that formal long-term contracts increase investment incentives but reduce firms' ability to respond to demand shocks, overall hurting allocative efficiency. Our paper is similar to Zahur (2022) in that we quantify the tradeoffs between the benefits of long-term relationships and their negative externalities on allocative efficiency, but differs in that we highlight an additional channel of such externalities. That is, the formation and high performance of long-term relationships increase spot market frictions by making the spot market thinner.

In financial settings, liquidity externalities—the fact that entry decisions by some participants increase the depth and liquidity of a market, thereby making that market more attractive to other participants—mirror the market-thickness externalities in the truckload setting. An implication of these liquidity externalities is that the prevalence of over-the-counter markets could be self-fulfilling while socially inefficient (Admati & Pfleiderer, 1988; Biais & Green, 2019; Pagano, 1989). Other arguments for the prevalence of over-the-counter markets include information asymmetry (Collin-Dufresne, Hoffmann, & Vogel, 2019; Lee & Wang, 2018), dealer heterogeneity (Dugast, Üslü, & Weill, 2019), and barriers to entry and insufficient competition on existing platforms (Allen & Wittwer, 2021). Our paper differs from these papers in the importance of match-specific gains from long-term relationships in our setting. In fact, we find that centralizing all transactions for truckload service results in substantial welfare loss,

precisely because of the complete loss of these match-specific gains⁶.

3 Institutional details

The US for-hire truckload freight industry offers an ideal setting to study the functioning of long-term relationships and their market impact. It is an economically important industry and one in which long-term relationships and spot arrangements coexist. The former are a central feature, with sophisticated institutions built around both the formation and management of relationships. This section provides institutional background and describes the nature of long-term relationships in this setting.

3.1 The US for-hire truckload freight industry

Trucking is the most important mode of transportation of US domestic freight. In 2019, trucks carried 64% of domestic shipments by weight and 72% of domestic shipments by value⁷. There are four main segments within the US trucking industry, separated by governance structure and size of shipments: truckload for-hire fleets, truckload private fleets, less-than-truckload, and parcel.

Our focus is on the for-hire truckload segment of the US freight trucking industry. 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 needs to hire a carrier (e.g., trucking company) for that service. This is in contrast to private fleets, which are vertically integrated carriers serving a single shipper⁸. In terms of shipment size, a shipment of a truckload carrier fills the entire truck. These carriers are concerned about reducing miles traveled empty, and thus unpaid, but not about how to optimally combine shipments to fill up their trucks. The latter is a key concern of less-than-truckload and parcel carriers. This means that truckload carriers face simpler routing decisions and rely less on economies of scale, a difference partially responsible for why the truckload segment is more fragmented than other segments (Ostria, 2003). The top 50 truckload fleets account for only about 10% of the seg-

⁶In the market for Canadian government bonds, Allen and Wittwer (2021) study the coexistence of a centralized platform and investor-dealer relationships. However, their analysis largely abstracts from match-specific gains from these relationships.

⁷These statistics are calculated using data from the Bureau of Transportation statistics.

⁸Such vertical contractual arrangements tend to be chosen by companies that prioritize quality and reliability of service, and typically have a dense network of truck movements that allow for efficient routing. In fact, shippers and carriers may maintain a portfolio of contractual arrangements (Acocella, Caplice, & Sheff, 2022a). For example, shippers with private fleets can sell excess capacity on their “backhauls”, and for-hire carriers can dedicate some of their capacity to some shippers.

ment's total revenue, and about 90% of truckload fleets have fewer than six trucks.⁹

Within the for-hire truckload segment of the US trucking industry, shipments are also separated by distance and trailer type. Long-hauls are shipments on lanes greater than 250 miles long. On such lanes, the ability of carriers to find backhauls is important (Hubbard, 2001). The common types of trailers are dry van, refrigerated, flatbed, and tanker. Our paper focuses exclusively on long-haul dry van truckload services.

Shippers and carriers in the US for-hire truckload industry engage in two main forms of transactions: long-term relationships and spot arrangements. Long-term relationships dominate this market, capturing 80% of total transacted volume; spot arrangements account for the remaining 20%.¹⁰

It is important to note that transactions on the spot market typically involve search and haggling. For example, shippers and carriers can post and search for available loads and trucks on electronic load boards. These load boards are marketplaces from which both sides can obtain contact information of potential matches, but rate negotiations are conducted offline.¹¹ Shippers and carriers can also be matched on digital matching platforms, which employ real-time matching and pricing, or via brokers. For our purpose, we treat all of these channels as a single spot market with search costs that potentially vary with spot market thickness.

3.2 Long-term relationships

There is an organized process that forms and manages long-term relationships in this setting, but contracts between shippers and carriers within their relationships are largely incomplete. The remarkably rich available data on how shippers and carriers form relationships and interact within relationships, together with the informal nature of these relationships, makes it an appealing empirical setting for our study.

Relationship formation. Long-term relationships are formed via procurement auctions. These auctions begin with shippers sending requests for proposals to different carriers, detailing their needs.¹² 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, which is typically one or two years. Contract rates are accompanied by a fuel program, typically proposed by shippers, that

⁹See <https://medium.com/@sambokher/segments-of-u-s-trucking-industry-d872b5fca913>.

¹⁰See <https://www.freightwaves.com/news/what-is-the-difference-between-trucking-contract-and-spot-rates>.

¹¹Figure 23 in Appendix E shows the search interface of DAT load board, the dominant load board for for-hire truckload service.

¹²Typically, shippers send out requests for proposals on multiple lanes simultaneously and carriers are free to bid on a subset of them.

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

Type	Order	Carrier	Rate (\$/mile)	Decision
Primary carrier	1	A	1.60	Reject
	2	B	1.44	Reject
Backup carriers	3	C	1.72	Accept
	4	D	1.89	

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

compensates carriers for changes in fuel costs.¹³ The shipper then chooses a primary carrier and a set of backup carriers in case transactions with the primary carrier do not materialize.¹⁴ Our analysis focuses on the relationships between shippers and their primary carriers.

Relationship management. Every interaction between a shipper and her carriers is automated and recorded by a Transportation Management System (TMS). When the shipper needs to transport a load on a lane, she inputs details of the load into her TMS, which automatically sends out offers to the carriers, sequentially in the order of their ranks, until one carrier accepts. Most carriers take less than one hour to respond to an offer.¹⁵ This process of sequential offerings is sometimes referred to as a “waterfall” process in other settings. Table 1 depicts an instance of this process.

While the auction determines an initial ranking of the carriers, the shipper can reorder this ranking at any point within the contract period. A “routing guide” keeps track of carriers’ updated ranks. The primary carrier is the top-ranked carrier, receiving most of the offers and typically accepting most of them.¹⁶ While backup carriers do not necessarily know their exact ranks, primary carriers know that they are top-ranked. This is because carriers need to plan ahead if they expect to service a large number of loads.

A key and unique feature of this setting is that contracts between shippers and carriers fix rates, but not volume. Shippers can influence the number of offers that each carrier receives

¹³The 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>.

¹⁴See Caplice (2007) for more details on this procurement process.

¹⁵The median response time is 41 minutes and 90% of all responses are within two hours. The full waterfall process typically takes less than three hours to complete.

¹⁶Sometimes, due to capacity constraints and other factors, offers are sent to the backup carriers first. See Appendix A for more details.

Figure 1: Spot market is the outside option of relationships

using their control over the routing guide, and carriers also face no legal recourse when rejecting loads. Such incompleteness in the contracts between shippers and carriers leaves room for potential opportunistic behaviors and for relational incentives to mitigate such behaviors.

Moral hazard and incentive schemes. In Harris and Nguyen (2021), we study the nature of the interactions between shippers and carriers within long-term relationships in this setting. Our two key findings are that (i) carriers can and do reject offers within relationships to take advantage of higher prices offered in the spot market, and (ii) shippers use a relational scheme to mitigate such temptation.

Figure 1 plots the movements of spot rates (in gray), average contract rates (dotted), and carriers' rejections (solid) over the period of our data sample.¹⁷ There is large temporal variation in spot rates. The market started soft with spot rates below contract rates, tightened over 2017 to reach its peak in 2018 with spot rates well above contract rates, and cooled down towards the end of the period. Average contract rates do adjust to spot rates, though only partially and with some lag. This means that there are some periods in which spot rates are much higher than contract rates. In such periods, carriers much more frequently decline shipment requests within relationships.

These data patterns suggest that the spot market is an important outside option for relationships. On the intensive margin, high premiums of spot rates over contract rates create a temptation for carriers to reject loads within long-term relationships. Shippers cannot observe if carriers truly do not have a truck available or if they are opportunistically declining to accept

¹⁷This figure is taken from Harris and Nguyen (2021), which has a more detailed discussion on the comovements of spot, contract, and rejection rates.

a higher priced offer on the spot market. Since shippers have imperfect monitoring of carriers' reasons for rejections, such a temptation constitutes a moral hazard problem. On the extensive margin, spot rates create an upward pressure on contract rates at the auction stage, affecting which relationships are formed and how relationship surplus is split between shippers and carriers.

To mitigate the moral hazard problem, shippers use their power to reorder the routing guide. That is, shippers can use the threat of demoting the current primary (top-ranked) carrier to a lower position on the routing guide to induce this carrier to accept more loads. In Harris and Nguyen (2021), we find that the shippers' incentive scheme takes a termination form: (i) higher rejection rates increase the likelihood of demotion, and (ii) once demoted, carriers hardly ever regain their primary status. Thus, we will model shippers' incentive scheme as a probabilistic termination strategy that conditions on carriers' past rejections. The estimated strength of this relational scheme will allow us to decompose the effects of relationships' intrinsic benefits from the effects of the dynamic incentives induced by such a scheme.¹⁸

4 Data

To capture the current market institution, we combine transaction-level data on long-term relationships and market-level data on spot arrangements. This section describes our data and provides empirical facts suggestive of the key tradeoff in our analysis: the match-specific gains from long-term relationships versus the efficiency of the spot market.

4.1 Transaction-level data on long-term relationships

We obtain detailed data on the interactions between shippers and carriers within long-term relationships from the TMS software provided by TMC, a division of C.H. Robinson.¹⁹ For each shipper and each lane of that shipper, we observe the details of all loads, including the origin, destination, distance, and activity date. Furthermore, we observe some aspects of shippers' input into the TMS software, including carriers' ranks and volume constraints, as well as information on the offers made to the carriers through the waterfall process, including their order, timestamps, contract rates, and carriers' decisions. We use these data to identify the primary carriers, when these primary carriers are replaced, and when auctions are held. We

¹⁸The shippers' incentive scheme is described in great details in Harris and Nguyen (2021, p. 30-35). Appendix D.1 presents our estimates of shippers' incentive scheme from a Probit specification, which will be used as an instrumental object in our empirical analysis. Consistent with Harris and Nguyen (2021), our estimates suggest that the shippers' incentive scheme is soft but generates dynamic incentives that are economically significant.

¹⁹C.H. Robinson is the largest third-party logistics firm in the United States.

define a relationship as the interactions between a shipper and a primary carrier on a lane, until that carrier is demoted and before the next auction.²⁰

Our data cover the period from September 2015 to August 2019. In total, we observe 1,186,413 loads and 2,367,704 offers between 54 shippers and 2020 carriers on 21,336 origin-destination pairs. These are all long hauls, with haul distance of at least 250 miles. We identify a total of 24,601 relationships, of which 13,171 are between a shipper and an asset-owner carrier. For our main analysis, we drop the relationships between shippers and brokers. The reason is that brokers act as intermediaries connecting shippers to carriers in the spot market; their cost structure and the nature of their relationships with shippers are thus very different from those of asset owners.²¹

Within the restricted data set, the average time between two auctions is 320 days, the average duration of a relationship is 33 offers, and the average number of offers is 7 loads per month. The average contract rate of primary carriers is \$1.82/mile, with a standard deviation of \$0.53/mile. In our sample, the premiums of spot rates over contract rates have mean \$0.04/mile, with a standard deviation of \$0.53/mile. On average, 70.3% of loads are accepted by primary carriers, 19.8% by backup carriers, and 8.9% of loads are fulfilled in the spot market. To simplify our analysis, we treat loads fulfilled by backup carriers as spot arrangements.

4.2 Market-level data on spot arrangements

We use spot rates data to capture the outside option of shippers and carriers in relationships and spot volume data to quantify the link between the thickness and efficiency of the spot market. These data come from DAT Freight and Analytics, the dominant freight marketplace platform in the US and the leading vendor of spot market data. DAT divides the US into 135 Key Market Areas (KMAs). We observe weekly summary statistics of spot rates and spot volume on each KMA-KMA lane. To merge these data with our transaction-level data on long-term relationships, we redefine origin-destination pairs in observed relationships at the KMA-KMA level. In total, there are 6,287 long-haul KMA-KMA lanes in our data on long-term relationships, out of 17,178 such lanes in the spot market data.

There are persistent differences in spot rates across lanes and large variation in spot rates over time. A regression of spot rates on lane fixed effects has an R^2 of 0.78; the average of these lane fixed-effects is \$1.58/mile, with a standard deviation of \$0.40/mile. Our analysis will control for time-invariant heterogeneity across lanes and exploit the large temporal variation

²⁰See Appendix A for details on how we construct indicators of primary status, demotion and auction events, and a graph of how these events are distributed over time in our sample period.

²¹See Appendix A, where we show that compared to asset-owner carriers, brokers have higher tendency to accept loads but are also more responsive to changes in spot rates.

of spot rates relative to contract rates to trace the value of relationships. The idea is that if a relationship has high value, the carrier's tendency to accept offered loads should be less sensitive to spot rates.

To proxy for spot market thickness on each lane, we use the average weekly number of loads on that lane that shippers post on DAT's marketplace. This marketplace is a "load board" where shippers post their demand and carriers search for available loads. There is large variation in spot market thickness, with 50% of the lanes having less than 20 spot loads per week and the top 1% of lanes having more than 500 spot loads per week. We exploit this variation to pin down the link between spot market thickness and efficiency.

4.3 Descriptive results

The magnitude of the two-way crowding-out effects between long-term relationships and the spot market depends on (i) the match-specific gains within relationships and (ii) the link between thickness and efficiency in the spot market. In one direction, the larger the intrinsic benefits of relationships, the weaker the crowding-out effect of the spot market. In the other direction, the stronger the link between spot market thickness and efficiency, the stronger the crowding-out effect of long-term relationships. This subsection provides evidence that relationships generate match-specific gains to participating parties but that they could also exert substantial negative externalities on the spot market by reducing spot market thickness.

Match-specificity. Patterns in the data suggest that match-specific gains are an important concern for both shippers and carriers.

On the shippers' side, two patterns in shippers' requests for shipment within relationships suggest that their gains from relationships are match-specific and potentially large. First, when selecting a primary carrier (the first carrier to receive load offers) in an auction, the shipper does not necessarily select the carrier that proposes the lowest contract rate. On the contrary, we see in Figure 2, which plots the distribution of the difference between the rates of the primary carrier and the lowest-rate backup carrier, that the primary carrier has the lowest contract rate in only two-thirds of the auctions. Among the remaining auctions, the median primary-backup price gap is 17 cents/mile. Such non-monotonicity in contract rates of shippers' ranking over carriers suggests that shippers also care about factors other than prices.

Second, shippers continue offering loads to their primary carriers even when lower-rate alternatives are available in the spot market. This pattern suggests that shippers' gains from transactions within relationships are potentially large. To show this, we run a regression of the variation in shippers' requests on variation in spot rates within the same contract period

Figure 2: Distribution of primary-backup price gaps

Table 2: Estimation results of Equation (1)

Sensitivity of shippers' offers to spot	
shipper 1	0.055 (0.011)
# Observations	197,749

(auction) on the same lane,²²

$$\ln \frac{\text{Volume}_{ia',\text{month}}^{\text{LT}}}{\text{Volume}_{ia'}^{\text{LT}}} = \beta_0^{\text{shipper}} + \beta_1^{\text{shipper}} \ln \frac{\tilde{p}_{ia',\text{month}}}{\text{Rate}_{ia'}} + \epsilon_{ia',\text{month}} \quad (1)$$

Table 2 shows that the estimate of β_1^{shipper} is negligible. This suggests that, unlike carriers, shippers do not defect from relationships when faced with spot temptation.

On the carriers' side, differences in the tendency of a carrier to accept offers in different relationships suggest that its gains from relationships are match-specific but potentially small. We consider the carrier with the largest number of relationships in our data set and approximate its acceptance tendency by a relationship-specific function of the normalized gap between spot and contract rates,²³

$$\Pr(d_{ij,t} = \text{accept}) = \text{Logit} \left(\beta_{ij,0}^{\text{carrier}} + \beta_{ij,1}^{\text{carrier}} \frac{\tilde{p}_t - p_{ij}}{\text{Std_Rate}} \right) \quad (2)$$

Table 3 presents the estimates of Equation (2) by mixed Logit. The likelihood ratio test of the mixed Logit specification against the pooled Logit specification has a Chi-square value of 1365.9, showing strong evidence of heterogeneity in the acceptance tendency of this single carrier across different relationships. Figure 3 demonstrates this heterogeneity in the distribution of the carrier's acceptance probabilities and sensitivities to spot rate. Here, acceptance probabilities are predicted for spot rates equal to contract rates; sensitivities to spot rates are measured by how much predicted acceptance probabilities decrease when spot rates increase

²² $\text{Volume}_{ia',\text{month}}^{\text{LT}}$ is the monthly number of requests that shipper i sends to her routing guide within the contract period of auction a on lane ℓ ; $\tilde{p}_{ia',\text{month}}$ is the median spot rate on lane ℓ in that month. We normalize these measures by their averages across all months of the contract period, $\text{Volume}_{ia'}^{\text{LT}}$ and $\text{Rate}_{ia'}$, to control for volume and rate differences across shippers and lanes.

²³ $d_{ij,t}$ is carrier j 's decision in period t in its relationship with shipper i on lane ℓ ; \tilde{p}_t is the spot rate on lane ℓ in period t ; p_{ij} is the contract rate; Std_Rate is the standard deviation of spot rates on lane ℓ .

Figure 3: Carriers' acceptance tendency

Table 3: Estimation results of Equation (2)

Sensitivity of a carrier's acceptances to spot				
carrier ij ⁰	Normal	0.46	1.84	0.28
carrier ij ¹		0.53	0.28	0.54
Heterogeneity		$\chi^2 = 1365.9$		
# Relationships		143		
# Observations		7,306		

Notes: Figure 3: The histograms are constructed by estimating Equation (2) separately for each relationship. The density curves are constructed by estimating Equation (2) with mixed Logit. Table 3: Mixed Logit results.

relative to contract rates by one standard deviation. In addition to showing large heterogeneity across relationships, the predicted sensitivities of this carrier's acceptance to spot rates suggest that its match-specific gains are potentially small. In two-thirds of its relationships, this carrier accepts less frequently as soon as spot rates exceed contract rates. In the median relationship, a one-standard-deviation increase in spot rates beyond this carrier's contract rate reduces its acceptance probability by 13 percentage points.

Spot market thickness and efficiency. At the market level, we find evidence suggesting a link between the thickness and the efficiency of 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 transactions taking place in the spot markets on lanes with higher potential for total market volume. We test this hypothesis by running the following regression on the difference in the growth rates of spot volume and long-term relationship volume as the total market volume increases,

$$\ln(\text{Volume}_{ss}^{\text{spot}}) - \ln(\text{Volume}_{ss}^{\text{LT}}) = \alpha_0 + \alpha_1 \ln(\text{Volume}_{ss}^{\text{total}}) + \text{controls} + \epsilon_{ss} \quad (3)$$

For this regression, lanes are defined at the state level. On each state-to-state lane, $\text{Volume}_{ss}^{\text{spot}}$ is the average weekly load posts in the spot market, $\text{Volume}_{ss}^{\text{LT}}$ is the average weekly loads accepted within long-term relationships, and $\text{Volume}_{ss}^{\text{total}}$ is the total for-hire truckload volume

in 2017 taken from the Commodity Flow Survey. We run a log-regression to avoid scaling issues between different data sets.²⁴ If $\beta_1 > 0$, the spot market takes a larger share of the total market volume as that total market grows, supporting our hypothesis.

An endogeneity concern in estimating this regression is that unobserved demand and supply factors could affect both total volume and the split between spot transactions and long-term relationships. First, demand patterns across industries may exhibit a correlation between total volume and preferences over forms of transactions. For example, shippers in some industries may have large total demand, but their lane-specific demand is infrequent and irregular; the latter prevents them from establishing long-term relationships. Second, unobserved cost factors may also make relationships easier or harder to establish. We mitigate the first concern by controlling for the frequency and consistency of load timing within observed relationships.²⁵ To address the second concern, we instrument for $\text{Volume}_{ss}^{\text{total}}$ with a demand shifter, the predicted trade flows between different states of the US from Caliendo, Parro, Rossi-Hansberg, and Sarte (2018). To construct these predicted flows, the authors first build a state-of-the-art trade model of the US economy that captures input-output linkages between different sectors, labor mobility, and heterogeneous productivities, but not the split between spot transactions and long-term relationships. They calibrate this model using 2012 data. The resulting predicted trade flows are for all modes of transportation, not just trucking.

Table 4 presents our estimates of Equation (3), showing in all specifications that spot volume increases faster than long-term relationship volume when there is greater demand for transportation service.²⁶ With regard to demand factors, the coefficient estimates of (OLS2) confirm that shipper- and lane-specific frequency and consistency of load timing make long-term relationships more desirable. However, the inclusion of these variables gives an estimate of β_1 similar to that of (OLS1); this suggests that demand factors, while important, do not create serious endogeneity concern in Equation (3). With regard to cost factors, the (IV) specification estimates a stronger link between spot market share and total market volume. This suggests that unobserved costs, which tend to reduce total market volume, may favor spot arrangements relative to long-term relationships.

To interpret the strength of the link between spot market share and total market thickness, we use the coefficient estimates in the (IV) specification to calibrate the shares of spot volume across all lanes. Figure 4 plots these shares against the total market volume. The fitted relationship shows that increasing the total market volume from 500 to 1000 loads per week

²⁴Relative to market-level data on long-term relationships, our microdata overrepresent the American Midwest. We control for this overrepresentation by including indicators of a lane's origin or destination being in the Midwest.

²⁵Ideally, we would also control for the lane-specific frequency and consistency of load timing of shippers who use spot arrangements. However, we do not have access to such data.

²⁶See Table 9 in Appendix D.2 for robustness checks.

Table 4: Estimation results of Equation 3

	ln(Volume ^{spot})		ln(Volume ^{LT})
	(OLS1)	(OLS2)	(IV)
ln(Volume ^{total})	0.285 (0.040)	0.269 (0.037)	0.345 (0.045)
ln(distance)	0.117 (0.080)	0.028 (0.075)	0.032 (0.077)
Frequency		0.156 (0.076)	0.159 (0.076)
Inconsistency		1.682 (0.319)	1.653 (0.321)
# Observations	588	588	588

Figure 4: Shares of spot market volume

Notes: Table 2: Frequency is the median average monthly volume on a lane; Inconsistency is the median coefficient of variation of loads in a week over the four weeks of a month. This regression aggregates spot and long-term relationship volumes to the state-to-state level and restricts to lanes with at least 10 relationships. Standard errors are in parentheses. Figure 3: The fitted line is constructed from $\hat{\alpha}_1 = 0.345$ from the IV specification. Two examples are included: Portland, Oregon to Syracuse, New York is a thin lane (200 loads/week); Buffalo, New York to Elizabeth, New Jersey is a thick lane (2000 loads/week). See Appendix C.3 for more details.

increases the share of the spot market from 10% to 20%. This finding suggests a potentially strong link between spot market thickness and the desirability of the spot market. Additionally, it is consistent with the finding in Hubbard (2001), that the share of spot relative to contractual arrangements in freight trucking increases with market thickness.

5 Model

To quantify the benefits of long-term relationships and their aggregate effects on the spot market, we need a model that captures: (i) the levels and heterogeneity of gains from long-term relationships, (ii) how these relationships interact with the spot market, and (iii) how spot market thickness is linked to spot market efficiency. To capture the match-specificity of relationships, we model the potential gains from the relationship between a shipper i and a primary carrier j on lane l as match-specific gains (g_{ij}^s, g_{ij}^c) to the shipper and carrier respectively over spot transactions. To allow for a potential link between spot market thickness and efficiency, we model per-load search cost for spot loads on lane l as a function of spot volume on that lane, $c_l = c_l(\text{Volume}^{\text{spot}})$. Long-term relationships and the spot market interact in two ways. First, the spot market serves as a clearing mechanism, fulfilling loads rejected within relationships. Second, the equilibrium volume split between long-term relationships and spot arrangements endogenously determines search costs on the spot market.

Figure 5: Model overview

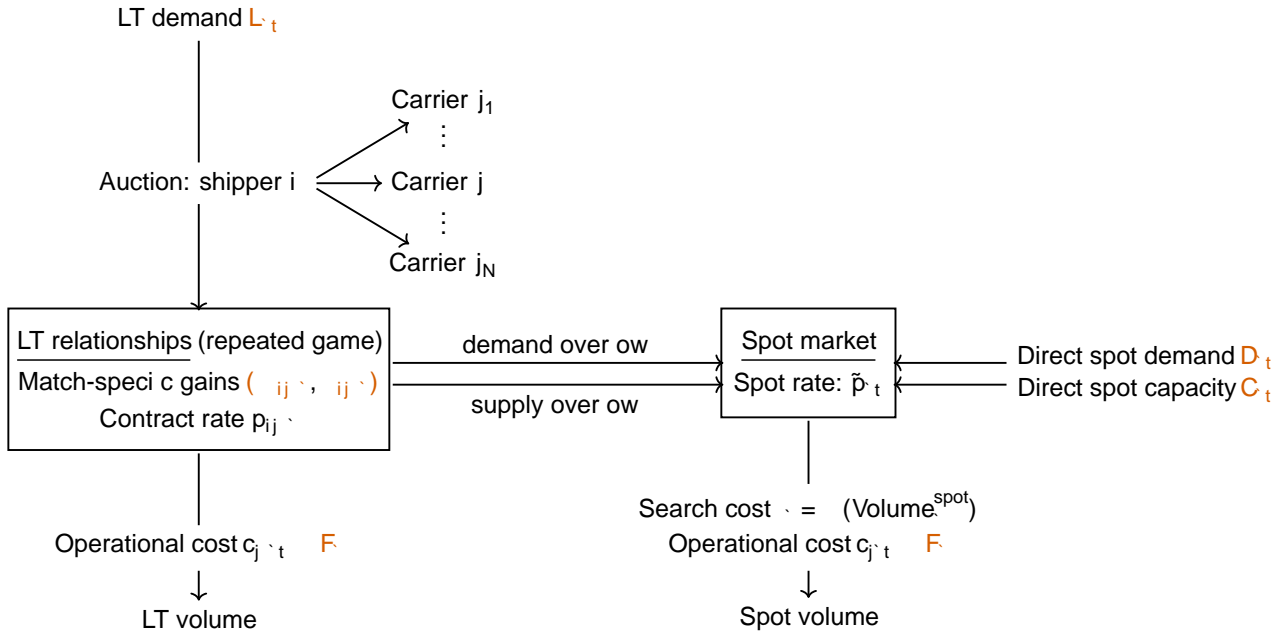


Figure 5 provides an overview of our model. Each long-term relationship goes through two stages. In the first stage, the shipper holds an auction to select a primary carrier. In the second stage, the shipper and primary carrier interact repeatedly under the fixed contract rate that was established by the auction. For each offer in the relationship, the carrier may reject because either spot rate or operational cost that period is high. This creates an overflow in both demand and supply from long-term relationships to the spot market. In equilibrium, such overflow, direct spot demand, spot capacity, search and operational costs pin down spot rate. We will use the observed equilibrium behaviors of individual shippers and carriers to recover key model primitives: the distribution of match-specific gains (g_{ij}^s, g_{ij}^c) , the distribution F of operational costs, and function ψ , which links search costs to spot market thickness. The market equilibrium condition will be used to recover the underlying supply and demand shocks, which are an input into our counterfactual analysis.

5.1 Timing and primitives of an individual relationship

We introduce the elements of an individual relationship in three layers: (i) shippers and carriers' per-period payoffs, (ii) the dynamics within a relationship, and (iii) the formation of that relationship.

Per-period payoffs. The relationship between a shipper i and a carrier j on lane l is characterized by a tuple $(g_{ij}^s, g_{ij}^c, p_{ij}^c, r)$ of relationship characteristics and a tuple (P, F, ψ) of

lane characteristics. Here, γ_{ij}^c is the match-specific gain of the shipper from transacting with the carrier on lane l ; γ_{ij}^c is the match-specific gain of the carrier from transacting with the shipper on lane l ; p_{ij}^c is the contract rate; β is the discount factor, reflecting the frequency of interactions; P^c is the spot process and F^c is the distribution of the carrier's operational costs on lane l ; $\kappa_l = \kappa(\text{Volume}^{\text{spot}})$ is the cost to a carrier of searching for a spot load on lane l , which depends on the average spot volume on that lane. Denote by \tilde{p}_t and $c_{j,t}$, respectively, the spot rate and the operational cost draw of the carrier on lane l in period t . The shipper's period- t payoff is $u_{i,t} = \gamma_{ij}^c - p_{ij}^c$ if she is served by the contracted carrier and $u_{i,t} = \tilde{p}_t$ if she is served by the spot market. The carrier receives the period- t payoff of $v_{j,t} = \gamma_{ij}^c + p_{ij}^c - c_{j,t}$ when delivering a load for the contracted shipper and $v_{j,t} = \tilde{p}_t - \kappa_l - c_{j,t}$ when serving the spot market. That is, a contracted load, if accepted, yields a premium (over a spot load) of γ_{ij}^c to the shipper and a premium of $\gamma_{ij}^c + \kappa_l$ to the carrier, including the carrier's savings on search costs.

Repeated game (dynamics within the relationship). In each period ($t \geq 1$) of a relationship between shipper i and carrier j on lane l , the shipper decides whether to terminate the relationship or offer a load to the carrier, and the carrier decides, if the shipper offers a load, whether to accept or reject it.

We allow the shipper to condition relationship termination on past decisions of the carrier. Let d_t denote the carrier's decision in period t : $d_t = \text{accept}$ if the carrier accepts the offered load; $d_t = \text{spot}$ if the carrier rejects the offered load to serve the spot market; $d_t = \text{idle}$ if the carrier rejects and remains idle. Denote by R_t an index summarizing carrier rejections in every period up to t . It is defined recursively by $I_R: (R_{t-1}, d_t) \rightarrow R_t = R_{t-1} + (1 - \alpha) \mathbb{1}\{d_t \neq \text{accept}\}$, where the weight α and the initial state R_0 are known. Let the spot rate follow an AR(1) process. The rejection index at the beginning of a period and the spot rate in the last period form a public (Markov) state $(R_{t-1}, \tilde{p}_{t-1})$. If the relationship terminates in some period, both the shipper and the carrier resort to the spot market for all future transactions; otherwise, the relationship continues to the next period starting at a new public state.

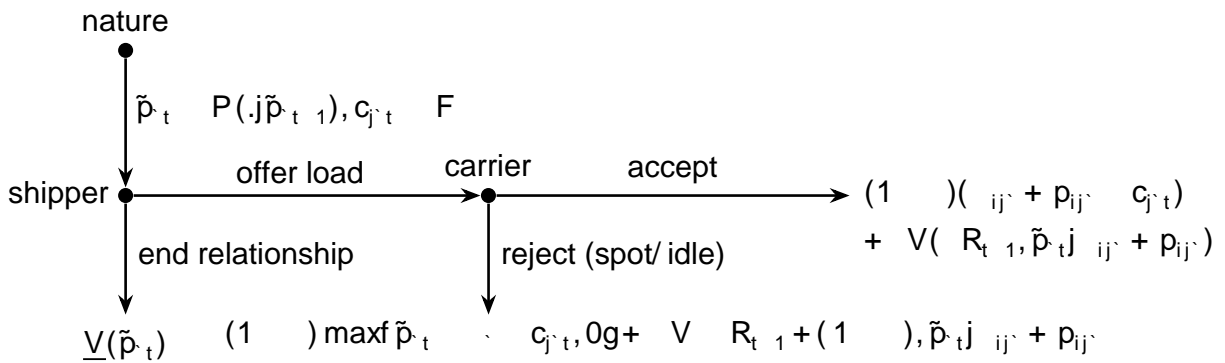
Assume that the shipper uses an incentive scheme $s_t: (R_{t-1}, \tilde{p}_{t-1}) \rightarrow [0, 1]$, which specifies for each rejection index and current spot rate the probability that the shipper maintains the carrier's primary status and offers it a load. Assume that the carrier strategy, $c_t: (R_{t-1}, \tilde{p}_{t-1}) \rightarrow \{\text{accept}, \text{spot}, \text{idle}\}$, specifies the carrier's optimal action for each rejection index and spot rate.

At the beginning of an auction, the seller announces an incentive scheme s_t that will apply to whomever wins the auction. Although it would seem restrictive to assume that the shipper does not condition the incentive scheme on the auction outcome, in our framework, s_t can be interpreted as the (average) incentive scheme perceived by all bidding carriers.

Assumption 1. The shipper's incentive scheme $s : (R_{t-1}, \tilde{p}_t) \in [0, 1]$ does not depend on the outcome of the auction. That is, s depends only on the characteristics of the lane and the discount factor δ .

Under Assumption 1, the expected payoff of the winning carrier j depends on the auction outcome only through its per-transaction rent $r_{ij} + p_{ij}$. Write $V(R_{t-1}, \tilde{p}_{t-1} | j, r_{ij} + p_{ij})$ for the expected payoff of the carrier in period t conditional on the rejection index R_{t-1} at the beginning of that period and the spot rate \tilde{p}_{t-1} last period, if this carrier's per-transaction rent is $r_{ij} + p_{ij}$. Write $\underline{V}(\tilde{p}_t)$ for the carrier's expected payoff from always going to the spot market, starting with \tilde{p}_t as the current spot rate. Figure 6 plots the timing of the stage game at state $(R_{t-1}, \tilde{p}_{t-1})$ and the carrier's payoffs in different outcomes.

Figure 6: The stage game at $(R_{t-1}, \tilde{p}_{t-1})$ and the carrier's discounted expected payoffs



In contrast, the expected payoff of shipper i depends on both her per-transaction rent $r_{ij} - p_{ij}$ and the carrier's per-transaction rent $r_{ij} + p_{ij}$, since the latter affects the carrier's tendency to accept loads within their relationship. This also means that the shipper might prefer a higher contract rate to a lower contract rate if the former induces significantly higher acceptance probability by the carrier, an idea similar to an efficiency wage. Thus, we write $U(R_{t-1}, \tilde{p}_{t-1} | j, r_{ij} - p_{ij}, r_{ij} + p_{ij})$ for the shipper's expected payoff in state $(R_{t-1}, \tilde{p}_{t-1})$ if the relationship induces per-transaction rents $(r_{ij} - p_{ij}, r_{ij} + p_{ij})$, and $\underline{U}(\tilde{p}_t)$ for the shipper's expected payoff from always going to the spot market given the current spot rate \tilde{p}_t .

Auction (formation of the relationships). At the auction stage ($t = 0$), a set of carriers propose contract rates and the shipper chooses a carrier with whom to form a relationship. We use subscripta to denote auction-specific variables. The timing of an auction is as follows:

- (i) Shipper i announces the expected frequency of interaction (λ_i) , other characteristics of the lane, and incentive scheme s .
- (ii) A set J_a of N carriers arrive.

- (iii) Pairs of shipper-carrier match-specific gains (v_{ij}, c_{ij}) are drawn i.i.d from the distribution G . The shipper's match-specific gain v_{ij} is observable to both her and carrier j , while the carrier's match-specific gain c_{ij} is privately known to carrier j alone.
- (iv) Each carrier j proposes a contract rate p_{ij} .
- (v) Shipper i chooses the carrier that maximizes her expected payoff as long as it is not lower than her expected payoff from the outside option of always going to the spot market.

A key assumption is that all carriers on the same lane have the same cost distribution and search costs, but they differ in the match-specific gains (v_{ij}, c_{ij}) that they will generate for each interaction with the shipper. Moreover, carriers' match-specific gains are privately known to carriers, whereas the shipper's match-specific gain with each carrier is known between the pair. In addition to providing tractability, these informational assumptions match certain features of the communication process between shippers and carriers. In her requests 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.

The informational assumptions above will allow us to transform carriers' bidding problem into the space of per-transaction rents. Match-specific gains affect shippers and carriers' expected payoffs only through these rents. First, since v_{ij} is known between shipper i and carrier j , by proposing a contract rate p_{ij} , the carrier essentially proposes a per-transaction rent $v_{ij} - p_{ij}$ to the shipper. That is, the proposed shipper's rent is the carrier's effective bid. Second, the shipper forms her expected payoff in each relationship from each carrier's effective bid and the carrier's rent; the latter is inferred in equilibrium. Finally, the assumption that carriers do not know the match-specific gains potentially generated by other carriers allows us to use empirical tools from the literature on independent private value auctions.

5.2 Equilibrium behaviors of individual shippers and carriers

This section derives three equilibrium conditions of individual relationships that will be used for identification. First, the winning carriers play optimally within their relationships with shippers. Second, given their expected payoff from a relationship, carriers bid optimally at the auction stage. Third, shippers select the relationships that yield them the highest expected payoffs, unless these payoffs are lower than what they would get from the spot market. Restricting our analysis to the class of symmetric monotone equilibria, we will write the last two conditions in the space of shippers and carriers' rents.

We start with the optimal dynamic play of the winning carrier j in the repeated game. Define the “full compensation” for this carrier by

$$\begin{aligned} & \bar{p}(R_{t-1}, \tilde{p}_{t-1}^{ij} + p_{ij}) \\ = & p_{ij} + \delta + \frac{1}{1-\delta} (V(R_{t-1}, \tilde{p}_{t-1}^{ij} + p_{ij}) - V(R_{t-1} + (1-\delta), \tilde{p}_{t-1}^{ij} + p_{ij})) \end{aligned} \quad (4)$$

and “transformed cost” by $\tilde{c}_{j,t} = c_{j,t} + \delta$. The optimal strategy σ_j of this carrier at each $(R_{t-1}, \tilde{p}_{t-1})$ depends on the relative ranking of $\bar{p}_{t-1} = \bar{p}(R_{t-1}, \tilde{p}_{t-1}^{ij} + p_{ij})$, \tilde{p}_{t-1} and $\tilde{c}_{j,t}$ as follows:

$$\sigma_j(R_{t-1}, \tilde{p}_{t-1}^{ij} + p_{ij}) = \begin{cases} \text{accept} & , \text{ if } \bar{p}_{t-1} \geq \max\{\tilde{p}_{t-1}, \tilde{c}_{j,t}\} \\ \text{spot} & , \text{ if } \tilde{p}_{t-1} > \max\{\bar{p}_{t-1}, \tilde{c}_{j,t}\} \\ \text{idle} & , \text{ if } \tilde{c}_{j,t} > \max\{\bar{p}_{t-1}, \tilde{p}_{t-1}\} \end{cases} \quad (5)$$

When a carrier decides whether to accept a load, it takes into account the contract rate p_{ij} , the match-specific gain \tilde{p}_{ij} , savings on search cost, and the effect of an acceptance on its continuation value. These components constitute the carrier's benefits from accepting a load within the relationship at each Markov state of the relationship, captured by the full compensation \bar{p} . Intuitively, when the carrier's per-period rent is higher, it is compensated more for an acceptance today both directly through higher compensation today and indirectly through higher compensation in the future.

Next, we focus on the class of symmetric monotone equilibria, which transform the bidding problem into the space of (per-transaction) rents.

Definition 1. (Symmetric monotone equilibria) An equilibrium of the auction and repeated game is a symmetric monotone equilibrium if there exists a strictly increasing and differentiable function $b : \mathbb{R}^+ \rightarrow \mathbb{R}$, referred to as the “effective bidding function”, such that

- (i) (Single indexing) The equilibrium (per-transaction) rents of shipper i and carrier j depend only on their total match quality $\tilde{p}_{ij} = \tilde{p}_{ij} + p_{ij}$,

$$\text{shipper } i\text{'s rent: } \tilde{p}_{ij} - p_{ij} = b(\tilde{p}_{ij}),$$

$$\text{carrier } j\text{'s rent: } \tilde{p}_{ij} + p_{ij} = \tilde{p}_{ij} + b(\tilde{p}_{ij}).$$

- (ii) (Monotone bidding) The shipper's rent $b(\tilde{p}_{ij})$ and the carrier's rent $\tilde{p}_{ij} + b(\tilde{p}_{ij}) - r(\tilde{p}_{ij})$ are both strictly increasing in \tilde{p}_{ij} .

- (iii) (Optimal selection) The shipper chooses the carrier j that maximizes her per-transaction rent subject to her expected payoff being no less than her outside option of always going

to the spot market,

$$j \in \mathcal{J}_a \text{ arg max}_{j \in \mathcal{J}_a} b(j) \quad \text{s.t. } U(R_0, \tilde{p}_0 | b(j), j) = b(j) \quad E[U(\tilde{p}_1) | \tilde{p}_0].$$

The first condition says that only the total match quality, rather than the relative magnitude of the match-specific gains v_{ij} and v_{ji} , affects which relationship is formed and how the gain from each transaction is split between the shipper and the carrier. That is, the observed price p_{ij} will adjust to reflect the relative magnitude of the shipper's versus the carrier's match-specific gains. The second condition implies that there is a one-to-one mapping between carriers' rents and effective bids, or shippers' rents. This condition is crucial to our identification argument; it will allow us to recover the distribution of shipper's rents from the distribution of carriers' rents. Finally, the third condition reduces a shipper's selection rule to choosing the carrier with the highest effective bid subject to her individual rationality constraint. The intuition is that in a symmetric monotone equilibrium, carriers that have higher effective bids are also those with higher rents, and thus would accept more frequently in a relationship. By choosing the carrier with the highest effective bid, a shipper thus maximizes both her per-transaction rent and the likelihood that such rent realizes. Appendix B.3 provides sufficient conditions for the existence of a symmetric monotone equilibrium.

Recall that G^* denotes the distribution of match-specific gains in random matches of shippers and carriers. Let $G^+ = G^*$ be the distribution of total match quality induced by G^* . Similarly, write G^{+P} for the distribution of carriers' rents and G^P for the distribution of shippers' rents induced by G^* and bidding function b . Key to our welfare analysis is the distributions of carriers and shippers' rents in relationships selected by the auction process. Denote these distributions by $[G^{+P}]^{1:N}$ and $[G^P]^{1:N}$ respectively. The following proposition summarizes all equilibrium conditions on shippers and carriers' behaviors that will be used for identification.

Lemma 1. (Equilibrium behaviors) Consider a symmetric monotone equilibrium with initial spot rate \tilde{p}_0 and effective bidding function b . The following hold:

- (i) (Optimal dynamic play) The winning carrier with per-transaction rent $v_{ij} + p_{ij}$ uses the optimal accept/reject strategy $\alpha_c(\cdot, j | v_{ij} + p_{ij})$ defined in Equation (5) when offered loads.
- (ii) (Optimal symmetric monotone bidding) For all type $j \in \mathcal{J}_a$,

$$b(j) = \arg \max_b G(b^{-1}(b))^{N-1} (V(R_0, \tilde{p}_0 | j, b) = b) \quad E[V(\tilde{p}_1) | \tilde{p}_0], \quad (6)$$

with $b(j)$ and $v_{ij} = b(j)$ both strictly increasing in j .

(iii) (Binding shipper's IR constraint) Let \underline{p} be the lowest match quality in a relationship. Then,

$$U(R_0, \tilde{p}_0 | b(\underline{p}), \underline{p}, b(\underline{p})) = E[U(\tilde{p}_1) | \tilde{p}_0]. \quad (7)$$

At every state and period of the repeated game, the carrier accepts only if its full compensation is higher than the compensation from the spot market and also higher than the sum of its operational and search costs. At the bidding stage, the bidding problem is as if carriers bid on the shipper's rent in a first-price auction with a reserved price determined by the shipper's individual rationality constraint. Notice that we do not impose optimality on the shipper's incentive scheme. One could in principle assume also that the shipper commits to an ex ante optimal incentive scheme or chooses a self-enforcing strategy, which is optimal period by period. Given that shipping is only a small component of shippers' business, we are less confident that such conditions would hold in reality, and did not want to impose it.

5.3 Market equilibrium condition

Next, we derive the market equilibrium condition that pins down the allocation of loads between long-term relationships and the spot market. Denote by L_t the measure of shippers who want to establish long-term relationships on lane ℓ in period t , by D_t the measure of direct spot demand by shippers who do not want to establish relationships and by C_t the spot capacity of carriers. Let $\nu_t(\cdot | \tilde{p}_t)$ denote the distribution over full compensations \bar{p} for potential relationships, conditional on the current spot rate. That is, $\nu_t(\cdot | \tilde{p}_t)$ captures both the extensive and intensive margins of volume in long-term relationships. Specifically on the extensive margin, $\nu_t(0 | \tilde{p}_t)$ measures potential relationships that are not formed and relationships that have ended. On the intensive margin, higher $\nu_t(\cdot | \tilde{p}_t)$ in the FOSD-sense means higher aggregate acceptance probability. Moreover, the status of each relationship as captured by the rejection index, is embedded in the measure $\nu_t(\cdot | \tilde{p}_t)$. The market equilibrium condition is

$$L_t + D_t = L_t \int_{\tilde{p}_t}^{\infty} F(\bar{p} | \cdot) d \nu_t(\bar{p} | \tilde{p}_t) + \left[\frac{L_t \nu_t(\tilde{p}_t | \tilde{p}_t) + C_t}{\int_{\tilde{p}_t}^{\infty} F(\bar{p} | \cdot) d \nu_t(\bar{p} | \tilde{p}_t)} \right] F(\tilde{p}_t | \cdot). \quad (8)$$

LT-relationship volume
spot volume

Within long-term relationships, loads are accepted if the full compensations \bar{p} of carriers are higher than the current spot rate and higher than the sums of carriers' search and operational costs. Loads offered but rejected in long-term relationships will be fulfilled in the spot market, either by carriers in relationships with low full compensations or by those not in relation-

ships.²⁷ A positive aggregate demand shock, either from an increase in demand for long-term relationships or from an increase in direct spot demand, increases the equilibrium spot rate.

For the welfare analysis of different market institutions, we keep fixed the long-term demand L_t , the short-term demand D_t , and the total capacity of carriers, including the spot capacity C_t and those that form relationships L_t . Additionally, we normalize the value of spot interactions for shippers who want to establish long-term relationships to zero. Thus, the market-level welfare of the current institution on lane l in period t is

$$W_t^0 = L_t \int_{\tilde{p}_t}^Z (E[j|\tilde{p}, \tilde{p}_t] - E[c_t|c_t, \tilde{p}]) F(\tilde{p}) d_{jt}(\tilde{p}|\tilde{p}_t) + [C_t + L_t \int_{\tilde{p}_t}^Z (\tilde{p}_t - j|\tilde{p}_t)] (E[c_t|c_t, \tilde{p}_t] - E[c_t|c_t, \tilde{p}_t]) F(\tilde{p}_t).$$

Alternative institutions that change the dynamics of long-term relationships modify the measure $d_{jt}(\tilde{p}|\tilde{p}_t)$ and the conditional expected match quality $E[j|\tilde{p}, \tilde{p}_t]$. For the extreme case with no long-term relationships, we set $d_{jt}(0|\tilde{p}_t) = 1$.

6 Identification and estimation

In this section, we discuss the identification of our model. We then specify parametric assumptions in an empirical model and explain our estimation procedure, which follows the steps in the identification argument closely.

6.1 Identification

Suppose that in each relationship we observe the contract rate p_{ij} , the duration of the relationship T_{ij} , and in each period $t \in T_{ij}$, the spot rate \tilde{p}_t and whether the carrier accepts or rejects, that is, whether $d_{jt} = \text{accept}$ or $d_{jt} = \text{spot, idleg}$. The observed relationships have unobserved match-specific gains $(\gamma_{ij}, \beta_{ij})$, lane-specific distribution F of operational costs, search cost s , and incentive scheme δ . Furthermore, suppose that we observe the number n_a of bidders in each auction who pass the shipper's individual rationality constraint and become either primary or backup carriers. For simplicity, assume that the discount factor is $\beta = \delta$. 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 spot process \tilde{p}_t is AR(1) and has $\text{supp}(\tilde{p}_t - j|\tilde{p}_{t-1}) = \mathbb{R}^+$ for every \tilde{p}_{t-1} .

²⁷For simplicity, we count backup carriers as spot carriers.

- (ii) The shipper's strategy satisfies that $\phi_s(R_{t-1}, \tilde{p}_t) > 0$ for all (R_{t-1}, \tilde{p}_t) .
- (iii) The underlying distribution of match-specific gains G^* has full support in \mathbb{R}^2 . Moreover, it induces an underlying distribution of match quality $G = G^* + \epsilon$ that has a strictly decreasing hazard rate, $g = G$.²⁸
- (iv) The cost distribution F is Normal(μ, σ).

Assumption 3. (Properties of full compensation schedules) Under the spot process and the shipper incentive scheme ϕ_s , the full compensation schedule $\tilde{p}(R_{t-1}, \tilde{p}_t | j + p)$ is:

- (i) strictly increasing in the carrier's rent $j + p$ for all (R_{t-1}, \tilde{p}_t) ,
- (ii) continuous in \tilde{p}_t for all R_{t-1} and $j + p$,
- (iii) bounded below by $j + p + \epsilon$ for all (R_{t-1}, \tilde{p}_t) .

Note that Assumption 3 is essentially an assumption on the underlying spot process and the shipper's incentive scheme. The most substantive assumption is (i).²⁹ Assumption (ii) holds under mild regularity conditions, and Assumption (iii) is satisfied if the incentive scheme ϕ_s is decreasing in R_{t-1} , that is, if the shipper punishes the carrier's rejections with a higher probability of demotion.

Proposition 1. (Full identification) Under Assumption 2 and Assumption 3, the model is fully identified within the class of symmetric monotone equilibria.

The shipper's incentive scheme ϕ_s is identified, since under Assumption 2, every Markov state (R_{t-1}, \tilde{p}_t) is observed.³⁰ The spot process P is identified from the realized path of spot rates. This section provides an identification argument for the following key primitives: the distribution F of operational costs, the search cost ϵ , and the joint distribution G^* of the match-specific gains of shippers and carriers on each lane.

These primitives are identified sequentially in four steps. First, we identify the distribution of carriers' operational and search costs. In this step, we exploit how carriers' tendency to accept offered loads varies across relationships and across lanes. The variation of such a tendency across different relationships on the same lane gives us different draws of the common cost distribution, and its variation across lanes with different spot market thickness pins down search costs. Given operational and search costs, the second step identifies the distribution of

²⁸For example, Normal distributions have strictly decreasing hazard rates.

²⁹The left panel of Figure 22 in Appendix E shows that this assumption is numerically verified under our estimates of the spot process and the shipper's incentive scheme.

³⁰Specifically, Assumption 2(i) ensures that every level of spot rate is observed, regardless of the current rejection state; (ii) ensures that every Markov state is non-absorbing; (iv) ensures a strictly positive probability of both acceptance and rejection in any Markov state.

carriers' rents from the distribution of observed acceptances. The idea is that, all else equal, carriers with higher rents accept more. Third, we recover the distribution of shippers' rents by exploiting the optimality of carriers' bidding and shippers' selection in a symmetric monotone equilibrium. This step uses objects recovered from the previous steps to construct the carriers' probability of and expected payoff from winning an auction, which are key inputs to the optimal bidding condition. Finally, we exploit observed prices to map the distribution of shippers' and carriers' rents to the joint distribution of match-specific gains.

6.1.1 Identification of the distribution of operational costs and search costs

The key observables in our identification of carriers' primitives are carriers' "acceptance schedules", or their tendency to accept a load at each Markov state of their relationship. Following from carriers' optimal dynamic play in Equation (5), a carrier's acceptance schedule depends only on its transformed rent $\tilde{r}_{ij} + p_{ij} + \dots$ and the distribution \tilde{F} of transformed costs, $\tilde{c}_t = c_t + \dots \sim \text{Normal}(\mu + \dots, \sigma^2)$.

Definition 2. (Acceptance schedules) Fix the shipper's strategy s , carrier's rent $\tilde{r}_{ij} + p_{ij}$ and lane characteristics (P, F, \dots) . Carrier j 's acceptance schedule is that carrier's tendency to accept a load at each Markov state (R_{t-1}, \tilde{p}_t) of its relationship with the shipper,

$$\Pr(d_t = \text{accept} | R_{t-1}, \tilde{p}_t) = 1 - \bar{p}(R_{t-1}, \tilde{p}_t) - \tilde{p}_t g \tilde{F}(\bar{p}(R_{t-1}, \tilde{p}_t)).$$

Notice that our assumption on the form of shippers' incentive schemes has two implications: (i) relationships evolve with a Markov state, and (ii) relationship terminations can occur on-path. The first implication means that what we observe about a carrier's dynamic play is precisely its tendency to accept loads at each Markov state of the relationship. Moreover, this tendency is fully observed in long-lasting relationships ($T \gg 1$). However, the second implication means that a significant proportion of relationships are relatively short-lived. Thus, we develop a two-step procedure. The first step takes advantage of long-lasting relationships to identify cost parameters. The second step identifies the distribution of carriers' rents, pooling relationships of all length. This section focuses on the first step.³¹

In this step, we separately identify cost parameters from carriers' match-specific gains by building on a thought experiment. Consider a hypothetical scenario when carriers only decide between "spot" and "idle". In this scenario, the probability that a carrier chose "spot" over "idle" at some spot rate would equal the CDF of the sum of this carrier's search and operational costs

³¹Note that long-lasting relationships are selected in match-specific gains or in cost draws. Our two-step procedure handles selection in match-specific gains. Appendix C.1.2 discusses how to correct for selection in cost draws in the first step.

Figure 7: Acceptance schedules and acceptance thresholds for a fixed rejection index

evaluated at this spot rate. Thus, the variation in the probability that a carrier chose “spot” over “idle” as spot rates vary would trace out the distribution of the sum of its search and operational costs. We build on this hypothetical scenario to pin down lane-specific cost parameters. In each relationship, we recreate this scenario locally by exploiting the observed level and sensitivity-to-spot-rate of the carrier's acceptance. Then, variation in carriers' acceptance across different relationships, with different contract rates or match-specific gains, on the same lane recreates the hypothetical scenario globally, pinning down cost parameters on this lane.

Figure 7 illustrates how the variation in the acceptance schedule induced by the variation in the level of carrier rent across relationships traces the distribution of transformed costs. The first two panels of Figure 7 plot the full compensation (black and blue lines) and optimal decisions of two carriers with different levels of carrier rent, fixing a rejection index and in the space of spot rate (x axis) and transformed cost (y axis). For each rent level, as spot rate increases, the full compensation starts from being higher than spot rate to eventually being lower than spot rate, 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 “idle” exactly at p . This means that the acceptance probability at p , which is the observed probability mass on the green vertical line, is the probability that the carrier chooses “spot” over “idle” at this level of spot rate were “accept” to not be an option, which is the unobserved probability mass on the red vertical line. That is, $(p, \tilde{F}(p))$ gives us one point on the distribution of transformed costs. Moreover, increasing the carrier rent shifts the acceptance schedule outwards. Such a shift results in a higher critical point, giving us another point on the distribution of transformed costs.

The last panel of Figure 7 illustrates how these critical points manifest as “jump” points in the observed acceptance schedules and are thus identified. At spot rates lower than the critical points, the carriers accept if their costs are below the full compensation, which happens with a positive probability. Acceptance probabilities jump to zero for any higher level of spot rate, because carriers would prefer spot to contracted loads even when cost draws are low. A caveat of this identification strategy is that it is empirically harder to pin down the left tail of the cost distribution, since “jump points” tend to be quite high, and more likely so in long-lasting relationships. However, acceptance schedules provide information about the cost distribution not only through their jump points, but also through the portion of these schedules to the left of the jump points. This argument provides intuition for Lemma 2 below. We delegate the formal definition of the “jump points” to Appendix B.1 and the identification proof to Appendix B.2.

Lemma 2. (Identification of the distribution of transformed costs) Fix lane characteristics. A carrier's acceptance schedule on lane l identifies at least one point on the distribution \tilde{F}_l of transformed costs, and the variation in carriers' rents identifies more points on \tilde{F}_l .

Lemma 3. (Identification of search costs and distribution of operational costs) If there is a demand shifter z independent of operational costs, then the variation in z and in the average spot volume, $\text{Volume}^{\text{spot}}$, identifies search costs $c_{j,t} = (\text{Volume}^{\text{spot}})$ and the cost distribution F , up to a constant.³²

Proof. By Lemma 2, the distribution \tilde{F}_l of transformed costs $\tilde{c}_{j,t} = c_{j,t} + (\text{Volume}^{\text{spot}})$ is identified. Since z is independent of $c_{j,t}$ and correlates with $\text{Volume}^{\text{spot}}$, it can be used as an instrument to non-parametrically identify $c_{j,t}$ in the range of spot volume. \square

6.1.2 Identification of the distribution of carriers' rents

Given the common cost distribution, we now identify the distribution of carrier rent. It is important to include all relationships in this step since excluding short-lived relationships would result in an upward bias of the recovered distribution of carriers' rents. To build intuition, consider the case where there are finitely many levels of carrier rent. The key idea is that the acceptance schedules associated with different levels of carrier rent are linearly independent, a property that ensures identification of finite mixtures of carrier rent.³³ Suppose that linear independence fails, that is, some acceptance schedule can be written as a linear combination of other acceptance schedules. Then none of these schedules can involve the highest

³²This constant has no bearing on our comparison of aggregate welfare between the current and alternative institutions. In the estimation, we pin down this constant by normalizing the median operational cost across lanes to industry estimates.

³³Kasahara 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.

acceptance schedule, since it is the only one in which a strictly positive probability of acceptance is observed near the highest “jump point”. Applying this argument iteratively from the highest to the lowest acceptance schedules yields a contradiction. Thus, the mixture of carriers' rents, if finite, is identified. Lemma 4 generalizes this intuition to a continuum of rent level. We present a direct proof of this lemma in Appendix B.2.

Lemma 4. (Identification of the distribution of carrier rent) Suppose that the shipper's incentive scheme s , search cost c , and the distribution F of operational costs are identified. Then the distribution $[G^{+p}]^{1:N}$ of winning carriers' rents is identified.

6.1.3 Identification of the distribution of shippers' rents

To identify the distribution $[G^{-p}]^{1:N}$ of shipper rent, we take advantage of the fact that in a symmetric monotone equilibrium, there is a one-to-one mapping between carrier rent ($r + p$) and shipper rent ($r - p$). To identify this mapping, we adapt the identification strategy in Guerre, Perrigne, and Vuong (2000). First, we transform the bidding problem into the space of carrier rent, the distribution of which is identified in the previous step. Second, we identify the monotone mapping between carriers' rents and their effective bids, or shippers' rents, from the first-order condition of optimal bidding in the space of carrier rent.

Lemma 5. (Identification of the distribution of shippers' rents) Suppose that the incentive scheme s , the distribution F of operational costs, search cost c , the distribution $[G^{+p}]^{1:N}$ of winning carriers' rents and the distribution of the number of bidders that pass the shipper's individual rationality constraint are identified. Then the distribution $[G^{-p}]^{1:N}$ of shipper rent is nonparametrically identified.

Proof. Consider an auction of shipper i with an initial spot rate \tilde{p}_0 . In a symmetric monotone equilibrium with an effective bidding function b and a rent function $r : \mathbb{R}^+ \rightarrow \mathbb{R}^+$, there exists a unique monotone mapping $b_r : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ defined by $b_r(r) = b(r - p)$. Thus, for carrier j with $\tilde{p}_j = \tilde{p}_0$, the optimal bidding condition reduces to choosing r that solves

$$\max_r [G^{+p}(r)]^{N-1} (V(R_0, \tilde{p}_0 | b_r(r)) - E[V(\tilde{p}_1) | \tilde{p}_0]).$$

The first-order condition gives

$$(N-1) \frac{g^{+p}(r)}{G^{+p}(r)} = \frac{\frac{\partial}{\partial r} V(R_0, \tilde{p}_0 | r)}{V(R_0, \tilde{p}_0 | r) - E[V(\tilde{p}_1) | \tilde{p}_0]} b_r^0(r). \quad (9)$$

Let \underline{r} denote the lowest level of carrier's rent in the support of $[G^{+p}]^{1:N}$, the shipper's IR

constraint can be rewritten as

$$U(R_0, \tilde{p}_0 | b(\underline{r}_\cdot), \underline{r}_\cdot) = E[U(\tilde{p}_\cdot | \tilde{p}_0)]. \quad (10)$$

Notice that the carrier's expected payoff as a function of carrier's rent r is identified from the incentive scheme s , cost distribution F and search cost ψ . To pin down b_r^0 from Equation (9), it remains to show that G^{+p} is identified on $[\underline{r}_\cdot, 1)$. We have for any rent level $r > \underline{r}_\cdot$, the distribution of winning carriers' rents satisfies

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

Moreover, the distribution of the number of effective bidders, who pass the shipper's individual rationality constraint, is Binomial $(N, 1 - G^{+p}(\underline{r}_\cdot))$. This means that $G^{+p}(\underline{r}_\cdot)$ is identified, which in turn identifies $G^{+p}(r)$ on $[\underline{r}_\cdot, 1)$ from the distribution of winning carriers' rents.

Finally, since the left hand side of Equation (10) is strictly increasing in its first argument, this equation pins down $b(\underline{r}_\cdot)$. Thus, b_r is identified, which in turn identifies the distribution of shippers' rents from the distribution of carriers' rents. \square

6.1.4 Identification of the distribution of match-specific gains

Finally, exploiting the fact that contract rates are observed, we identify the distribution of carriers' rents conditional on contract rates and thus, the joint distribution of shippers and carriers' match-specific gains.

Lemma 6. (Identification of the distribution of match-specific gains) Given the shipper's punishment scheme s , the common cost distribution F , search cost ψ , the number of effective bidders and the contract rates in all relationships, the distribution $[G^{+pj}]^{1:N}$ of winning carriers' rents conditional on contract rates is identified. It follows that the joint distribution $[G^\cdot]^{1:N}$ of shippers and carriers' match-specific gains is identified.

Proof. That $[G^{+pj}]^{1:N}$ is identified is an extension of Lemma 4 by conditioning on observed contract rates. Since the distribution $[G_p]^{1:N}$ of contract rates is observed, it follows that the joint distribution $[G^{+pj}]^{1:N}$ of carriers' match-specific gains and contract rates is identified. Furthermore, notice that Lemma 5 identifies the monotone equilibrium mapping b_r from carrier rent to shipper rent, that is, $b_r(r + p) = p$. This establishes a one-to-one mapping from (r, p) to $(b_r(r + p), p)$, completing the proof that $[G^\cdot]^{1:N}$ is identified. This means that the fundamental distribution of match-specific gains G^\cdot is identified but only for $r + p \leq \underline{r}_\cdot$, since we do not observe potential relationships that fail shippers' IR constraints. \square

6.2 Empirical model

Three features of the data require adaptations of our model. First, different shippers and carriers, through negotiations in the spot market, settle on different spot rates, but we only observe summary statistics of spot rates. To address this, our empirical model allows for idiosyncratic noise in the spot rate observed by a carrier at the time it decides whether to accept a load. Second, there is large persistent heterogeneity across lanes, which need to be controlled for in a pooled estimation. We use three variables to control for such heterogeneity: $\text{Rate}^{\text{spot}}$, $\text{Volume}^{\text{spot}}$, and Distance , which are respectively the average spot rate across time, the average spot volume across time and the average distance of a trip on a KMA-KMA lane.³⁴ For ease of interpretation, our estimation treats match-specific gains, rates and costs on a per-mile basis. Third, the truckload freight market goes through phases, as macroeconomic conditions change. A “soft” (“tight”) market is one in which demand for truckload service is lower (higher) than truckload capacity. Our empirical model takes the unit of a relationship to be between a shipper and a carrier on a lane in an auction period. We assume that the market phase can differ across auctions, but does not change within auctions.³⁵

Spot variance and spot process. Let \tilde{p}_t denote the mean spot rate on lane l in period t . Assume that at the time of decision making, carrier j faces spot rate $\tilde{p}_t + \epsilon_{jt}$, where $\epsilon_{jt} \sim \text{Normal}(0, \sigma^2)$. Assume further that future mean spot rates depend on the current mean spot rate through the following AR(1) process,

$$\frac{\tilde{p}_t}{\text{Rate}} = \alpha_0 + \alpha_1 \frac{\tilde{p}_{t-1}}{\text{Rate}} + \epsilon_t,$$

where t denotes calendar day. The spot process as perceived in a relationship adjusts this calendar-based spot process to the frequency of the shipper and carrier's interactions. Under these two assumptions, the Markov state of a relationship includes the rejection index and mean spot rate, (R_{t-1}, \tilde{p}_t) , and the observed acceptance schedule is “smoothed” out,

$$\Pr(d_t = A_j | R_{t-1}, \tilde{p}_t) = \frac{\bar{p}(R_{t-1}, \tilde{p}_t) \tilde{p}_t}{\int \bar{p}(R_{t-1}, \tilde{p}_t) \tilde{p}_t \tilde{F}_{ia}(\bar{p}(R_{t-1}, \tilde{p}_t))},$$

³⁴We take Distance to be the practical mileage on a KMA-KMA lane, which is the industry's estimate of the most likely distance of a trip from a KMA-origin to a KMA-destination.

³⁵Acocella, Caplice, and Shef (2020) identify breaks in the time series of rejection rates to define market phases. They identify the period from April 15, 2016 to April 14, 2017 as a full year in a soft market, and the period from October 1, 2017 to September 30, 2018 as full year a tight market. Since we assume that each relationship belongs to a market phase, we use the time of the auction to classify a relationship's market phase. Specifically, relationships that start in 2017 or 2018 belong to a tight market.

where \tilde{F}_{ia} is the auction- and lane-specific distribution of transformed costs.

Shippers' strategies. We specify a Probit model for demotion probability,

$$1 - s(R_{t-1}, \tilde{p}_t) = (\beta_0 + \beta_1 R_{t-1} + \beta_2 X_{i,t} + \beta_3 R_{t-1} X_{i,t}), \quad (11)$$

where $X_{i,t}$ is a tuple including the normalized spot rate $\tilde{p}_t = \text{Rate}$, the log of average monthly volume and the average coefficient of variation of weekly volume on lane ℓ of shipper i across all months of the auction period. As argued in Harris and Nguyen (2021), the latter two variables affect the desirability of a shipper-lane: (i) the frequency of interactions affects the continuation values of the relationship of both the shipper and carrier, and (ii) the consistency of offers captures the extent to which the carrier can benefit from planning.

Operational and search costs. Let (per-mile) operational costs of carrier j on lane ℓ of shipper i in auction a be distributed as a Normal distribution with mean c_{ia}^c and common standard deviation σ^c ,

$$c_{j,t} \sim \text{Normal}(c_{ia}^c, \sigma^c).$$

Thus, the transformed cost is distributed as $\text{Normal}(\tilde{c}_{ia}^c, \sigma^c)$, where $\tilde{c}_{ia}^c = c_{ia}^c + \beta$.

To decompose the transformed cost into operational and search costs, we make two further parametric assumptions. First, per-load search costs are linked to spot market thickness through a scale efficiency parameter β_1 , giving per-mile search costs³⁶

$$s = \frac{\beta_0 + \beta_1 \ln(\text{Volume}^{\text{spot}})}{\text{Distance}}. \quad (12)$$

Second, we allow operational costs to vary flexibly with distance, differ across market phases (soft or tight), and have unobserved differences. Specifically,

$$c_{ia}^c = \beta_2 1_{\text{tight}}^a + h(\text{Distance}) + \epsilon^c + \eta_{ia}^c, \quad (13)$$

where h is a flexible function. The regression of the mean of transformed costs $\tilde{c}_{ia}^c = c_{ia}^c + \beta$ as defined by Equations (12) and (13) has a potential endogeneity issue: lanes with higher unobserved cost shifter ϵ^c tend to have lower volume in equilibrium. We resolve this issue by instrumenting for realized spot volume with a demand shifter, the predicted trade flows across

³⁶Our functional-form assumption on per-mile search costs comes from our interpretation of search costs as fixed costs. This interpretation implies that per-mile search costs and per-mile thickness externalities are smaller on longer lanes. We empirically verify this implication by considering alternative functional-form assumptions on search costs. See Appendix 10 for our estimates of search costs under different specifications.

US states (Caliendo, Parro, Rossi-Hansberg, & Sarte, 2018).

Number of bidders. We allow for a stochastic number of bidders in each auction, N_a Binomial(N, q). When bidding, a carrier only knows (N, q) and not N_a .

6.3 Estimation procedure

Our estimation procedure follows the steps of the identification argument closely. First, we obtain instrumental objects, including the discount factor, the distribution of the number of bidders, the spot process, and the shippers' incentive scheme. These objects are input into the remaining steps, in which we estimate the cost parameters, the distribution of rents and the distribution of match-specific gains.

Using relationships of at least 50 offers, we estimate cost parameters in two steps. First, we estimate parameters of transformed costs by maximum likelihood (Rust, 1994). The likelihood contribution of each relationship is the likelihood of the carrier's accept/reject decisions. This likelihood contribution depends on the auction-specific mean \tilde{c}_{ia} of transformed costs, the common cost variance σ_c^2 , and the carrier's transformed rent $(\rho_{ij} + p_{ij} + \gamma)_{ij}$. For each relationship and each set of parameter values, we use a fixed-point algorithm to solve for the carrier's value function and optimal strategy. The full likelihood aggregates all relationships and is maximized in two layers. In the outer loop, we search for the value of the common cost variance σ_c^2 on a grid. In the inner loop, we jointly estimate auction-specific means $(\tilde{c}_{ia})_{ia}$ of transformed costs, and relationship-specific carriers' transformed rents $(\rho_{ij} + p_{ij} + \gamma)_{ij}$. Second, we decompose the estimated means of transformed costs into search and operational costs by estimating the following equation by two-stage least squares,

$$\tilde{c}_{ia} = \frac{\gamma_1 \ln(\text{Volume}^{\text{spot}})}{\text{Distance}} + \gamma_2 1_{\text{tight}}^a + h(\text{Distance}) + \sigma_c^2 + c_{ia}^c, \quad (14)$$

using predicted trade flows (Caliendo, Parro, Rossi-Hansberg, & Sarte, 2018) as an instrument for spot volume. Since we use a flexible function h to control for distance, this regression does not pin down the base per-load search cost γ_0 . To obtain an estimate of this parameter, we normalize the median operational costs across all lanes to industry estimates³⁷.

Given our cost estimates, we estimate the distribution of rents and the distribution of match-specific gains. We perform this step separately for each of ten clusters of lanes, to ensure

³⁷Using 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.

a sufficient number of relationships in each lane-specific cluster, and separately for soft and tight market, to allow for different demand and supply factors across the two market phases. Lane-specific clusters are constructed by K-means clustering based on the average spot rate, average spot volume, and distance. The idea is to identify the latent groups of lanes, as defined by fundamental demand and supply factors, by clustering on informative equilibrium fixed effects.³⁸ Thus, relationships in each of our clusters are potentially different in match-specific gains but similar in other characteristics.

Within each cluster, the distribution of shippers and carriers' rents, and their match-specific gains are estimated in three steps. First, we use an EM-algorithm (Train, 2008) to estimate the distribution of winning carriers' rents. Second, we piece together the monotone mapping between carriers' rents and shippers' rents from the first-order condition of carriers' bidding evaluated at the percentiles of the distribution of winning carriers' rents. For the initial condition of the bidding function, we take the fifth percentile of this distribution as the lowest carrier rent and pin down the lowest shipper rent via shippers' individual rationality condition.³⁹ Finally, to obtain the fundamental distribution of match-specific gains, we use simulated methods of moments to match the distribution of carrier rent conditional on different bins of contract rates. See Appendix C for a detailed description of our estimation procedure.

7 Estimates of match-specific gains, search costs, and operational costs

This section presents our estimates of key model primitives. These estimates suggest that long-term relationships generate large expected surplus to participating parties, but they exert substantial negative externalities on the spot market. We find that on a typical lane, doubling the thickness of the spot market reduces search costs by an amount equivalent to reducing operational costs by 29%. High expected surplus from relationships come from high match-specific gains, and the fact that the current fixed rate contracts and relational schemes perform fairly well at realizing these gains. We estimate that the median relationship achieves 44% of the first best surplus for individual relationships, with significantly better performance for relationships with higher match-specific gains.

Our estimates of instrumental objects, including the discount factor, number of bidders and relational scheme can be found in Appendix D.1. Consistent with Harris and Nguyen (2021),

³⁸See Bonhomme and Manresa (2015), who propose clustering methods to identify latent groups based on observed patterns of heterogeneity in panel data. See Bester and Hansen (2016), who study grouped effects estimators.

³⁹Note that the lower is the carrier's rent, the higher is the shipper's rent required for the relationship to pass the shipper's individual rationality constraint. Moreover, the tail of the estimated distribution of carriers' rents tend to have larger errors. Thus, our choice of the fifth percentile of the distribution of carriers' rents as the lowest level of carrier rent errs on the side of not including estimated shipper rents.

Table 5: Estimates of cost determinants

	Estimate	95% CI
Scale efficiency (α_1)	255.85	(-402.97, 146.26)
Tight market	0.54	(0.39, 0.64)

Note: Standard errors are constructed by bootstrapping at the auction level.

Figure 8: Estimated distribution of search and operational costs across lanes

we find that shippers punish carriers' rejections by increasing the probability of demotion in future periods. While soft, shippers' incentive scheme generates dynamic incentives that are economically significant.

7.1 Operational and search costs

One noteworthy finding on search costs is that increasing the thickness of the spot market reduces search costs. This supports the idea hypothesized by Kranton (1996), that the formation of long-term relationships can crowd out the spot market, making it thinner and less efficient. Table 5 reports the estimates of the parameter for scale efficiency (α_1) and the average difference in operational costs between the tight and soft market. Our estimate of α_1 is negative and economically significant. Specifically, we estimate that doubling the spot volume on a median lane of 500 miles decreases search costs by \$0.35/mile, an amount equivalent to 29% reduction in operational costs.⁴⁰

Figure 8 plots the distribution of estimated search costs and operational costs across lanes. The variation in estimated search costs comes from the large variation in spot volumes across lanes, and the residual variation in costs is captured by operational costs. To decompose search and operational costs from their sum, we relied on our estimate of α_1 and a normalization.

⁴⁰Table 10 in Appendix D.2 presents different specifications of (per-mile) search costs. The results show a strong link between search costs and spot volumes and that this link is weaker on longer lanes.

Figure 9: Distribution of match-specific gains and match quality

Note: The dashed lines on the left panel indicate the median levels of carrier's match-specific gain (including savings on search costs) and shipper's match-specific gain in realized relationships.

Specifically, we calibrated the base per-load search costs (\hat{c}_0 in Equation (12)) to match the median of our estimated operational costs to the industry estimate.⁴¹ We estimate that the median search cost is \$0.35/mile, equal to 29% of the median operational costs. Note that while these estimates are sensitive to our normalization method, they are not crucial to our welfare comparison, which relies mostly on our estimate of the scale efficiency $\hat{\alpha}_1$.

Finally, we find that operational costs are \$0.54 / mile higher in a tight market than in a soft market. This reflects the capacity crunch reported during the period from early 2017 to the end of 2018.⁴² This also means that the increase in rejection rates observed in Figure 1 is due to both an incentive effect and an aggregate cost effect.

7.2 Match-specific gains

We find large and heterogeneous match quality in realized relationships, which is accounted for mostly by shippers' match-specific gains. The left panel of Figure 9 plots the density of the joint distribution of carriers' match-specific gains including savings on search costs ($\hat{g} + \hat{c}$)

⁴¹Normalizing the median operational cost across all lanes to the industry estimate pins down $\hat{c}_0 = 1494.068$. In addition, the standard deviation of operational costs is estimated to be $\hat{\sigma}_c = 1.2$ (\$/ mile).

⁴²Industry reports suggest that the capacity crunch during 2017 and 2018 is driven by both rising demand for trucking services and supply-side factors, such as the Electronic Logging Device mandate and the enforcement of Hours of Service Law starting in January 1st, 2018. This capacity crunch is also reflected in the hourly wage of truck drivers. According to data from the American Transportation Research Institute, average hourly wage of truck drivers in the period from 2008 to 2020 peaked in 2018.

and shippers' match-specific gains (π) in realized relationships. Darker colors demonstrate values of shippers and carriers' match-specific gains with higher density. We find that in the median relationship, the shipper's match-specific gain is \$1.02/mile, accounting for 89% of the sum of the shipper and carrier's match-specific gains, or their match quality. The median carrier's match-specific gain is relatively small (\$0.12/mile) and in fact smaller than the median savings on search costs (\$0.35/mile). These findings are consistent with our descriptive results in Section 4.3, that (i) shippers do not reduce volume offered within relationships when spot rates are low, while (ii) a fair share of carriers accept much less frequently when spot rates are high. Moreover, the slight negative correlation between shippers and carriers' match-specific gains in realized relationships is due to the selection of relationships with high match quality.

The right panel of Figure 9 plots the quantile distribution of match quality, including and excluding savings on search costs. Even when savings on search costs are excluded, the estimated match quality is large and heterogeneous across relationships, increasing from \$0.42/mile to \$1.35/mile from the 25th to the 75th percentiles. Such large heterogeneity in match-quality reflects the importance of auctions in facilitating the formation of high-value matches in the truckload setting. Moreover, excluding savings on search costs, 10% of relationships have negative intrinsic match quality. This means that the dominance of long-term relationships in the current institution is self-fulfilling only to a limited extent.

7.3 Shippers and carriers' rents

The shipper and carrier's match-specific gains affect welfare only through how their match quality ($\pi + p$) is split into the carrier's per-transaction rent ($\pi + p$) and the shipper's per-transaction rent ($\pi - p$) by the equilibrium contract rate (p). Since these rents are realized only when the carrier accepts an offer of the shipper, the role of the equilibrium contract rate is not merely distributional. When the carrier has a higher rent from acceptance, it accepts more frequently, realizing the match-specific gains for both itself and the shipper.

The equilibrium split of total match quality between shippers and carriers is driven by three forces. First is individual rationality: since carriers have the option to reject shippers' offers as spot rates vary, shippers require large average rents to benefit from relationships, while carriers can benefit from relationships even with small average rents. Second is a competition effect: that carriers bid for relationships in auctions makes the split of match quality into rents more favorable towards shippers. Third is an "efficiency wage" effect: shippers may prefer leaving higher rents to carriers to induce more acceptances.

We estimate large and heterogeneous rents for shippers and carriers from long-term relationships. Figure 10 plots the quantile distribution of average carriers and shippers' rents

Figure 10: Quantile distribution of per-transaction normalized rents

Note: Carrier's normalized rent: $\frac{(\text{Rate} - p)}{\text{Rate}}$; shipper's normalized rent: $\frac{p}{\text{Rate}}$; where Rate is the mean spot rate over the entire sample period. The normalization is performed separately for each cluster of lanes, and the plots are of averaged normalized rents.

normalized by the mean spot rates, across soft (blue) and tight (red) markets. In the median relationship, the carrier has a normalized rent of 10% and the shipper has a normalized rent of 58%. From the 25th to the 75th percentiles, the carrier's normalized rent increases from 5% to 39% and the shipper's normalized rent increases from 45% to 65%.

Finally, carriers are compensated more in a tight market than in a soft market. The reason is that the better outside option of carriers in a tight market, when spot rates are high, improves their outcomes in auctions for relationships due to both the competition and efficiency wage effects. However, the increase of 10% in carriers' normalized rents does not fully compensate carriers for the increase in spot rates from a soft to a tight market.

7.4 Welfare analysis of individual relationships

We transform our estimates of the incentive scheme and per-transaction rents into shippers and carriers' expected surplus from long-term relationships over spot transactions. That is, this exercise takes into account both the intensive margin (equilibrium rejections) and extensive margin (equilibrium demotions) of relationships. We then benchmark the joint expected surplus of an individual relationship in the current institution against the relationship-level first-best surplus. The latter requires that (i) relationships never end, and (ii) the first-best outcome is achieved in each period. Specifically, for a relationship with match quality (inclu-

Figure 11: Expected surplus from long-term relationships

Note: The dash vertical lines indicate the median match quality (including savings on search cost) in each market.

sive of savings on search costs) $+ + 0$, the carrier should never service the spot market; it should accept when its cost after internalizing the joint match-specific gains is less than the shipper's payment in the spot market, $c_t (+) \tilde{p}_t$, and remain idle otherwise. Figure 11 plots the expected surplus from long-term relationships (blue line) under current fixed-rate contracts, how it is split between shippers (red area) and carriers (blue area), against the relationship-level first-best surplus (green line).

One of our main findings is that the current fixed-rate contracts and incentive scheme do a fair job at capturing the first-best surplus, with significantly better performance in relationships with higher match quality. Specifically, the median relationship achieves 44% of the relationship-level first-best surplus; this figure increases from 27% to 62% from the 25th to the 75th percentile of match quality. The reason for such heterogeneity is that relationships with lower match quality face a more serious moral hazard problem and have less room to use contract rates as an incentive instrument.

In terms of distributional effects, we find that shippers have a larger share of total expected surplus. At the median relationship, shippers enjoy 75% of total expected surplus. This reflects our estimate that, on average, there are three effective bidders per auction. The share of carriers in total expected surplus slightly increases with their relationships' match quality. This reflects the higher information rents of carriers with higher match quality.

7.5 Key insights

Our findings show that the two-way crowding-out effects between long-term relationships and the spot market, as hypothesized by Kranton (1996), are large in our setting. On the one hand, long-term relationships result in a thinner spot market with significantly higher search costs. Specifically, were we to double the thickness of the spot market, search costs would reduce by about \$0.35/mile. On the other hand, the current fixed-rate contracts allow the spot market to crowd out long-term relationships, achieving 44% of the relationship-level first-best surplus in the median relationship. Furthermore, the crowding-out effect of the spot market is stronger in relationships with a lower match quality. Such selectivity tends to be beneficial to market-level welfare. This is because to achieve the same level of spot market thickness, it is generally more efficient to forgo transactions that generate low match quality.

These findings beg the following questions: Is the current “high-relationship” equilibrium socially optimal or just self-fulfilling? What are the welfare effects of increasing the spot market thickness or improving the performance of long-term relationships? The next section will shed light on these questions.

8 Market-level welfare under alternative institutions

At the market level, there is a tradeoff between realizing more match-specific gains generated by long-term relationships and maintaining greater thickness of the spot market. We quantify this tradeoff by comparing the market-level welfare of the current institution to alternative institutions that change the share or performance of long-term relationships. The first counterfactual institution is a centralized spot market for maximal spot market thickness. The second counterfactual replaces all fixed-rate contracts with the individually optimal index-priced contracts. We find that a centralized spot market would result in substantial welfare loss, and index pricing would be beneficial, though only in periods with high demand. We also construct an upper bound on the market-level first-best surplus and find that the current institution achieves 40% of this surplus on medium trips of 500 miles and 60% of this surplus on long trips of 1000 miles.

8.1 Economic tradeoffs

The welfare effects of different market institutions depend on the own benefits of long-term relationships and the spot market, as well as how these two forms of transactions interact. First, long-term relationships generate large match-specific gains to participating parties, while the spot market, by centralizing more transactions, may reduce search costs and improve allocative

cost efficiency. Second, there are two-way crowding effects between relationships and the spot market in the US for-hire truckload freight industry. On the one hand, the spot market creates carriers' moral hazard problem, crowding out low-value relationships. On the other hand, the formation of long-term relationships results in higher search costs in a thinner spot market. We quantify the welfare effects of these economic forces in two counterfactual exercises.

A centralized spot market (no relationships). The first counterfactual is a spot platform that centralizes all transactions in the market. This achieves the optimal scale efficiency, thus reducing search cost on the spot market and increases allocative cost efficiency; the tradeoff is the complete loss of match-specific gains in long-term relationships.

Index-priced contracts (optimal relationships). This counterfactual keeps the auctions for relationship formation but replaces all fixed-rate contracts with the individually optimal index-priced contracts. These contracts take the standard idea in the contract literature that to solve a moral hazard problem in a principal-agent relationship, the principal should “sell the firm” to the agent. And to screen out the best relationship, the principal should ask all agents to bid on the contract. In an index-priced contract, it means that the shipper transfers all of the rents from each realized transaction to the carrier, and asks all carriers to bid on a fixed fee to be paid for each offer, regardless of whether the carrier accepts or rejects. While such index-priced contracts eliminate the moral hazard problem, achieving the first-best surplus of individual relationships, they would exacerbate the negative externalities of relationships on the spot market.

Definition 3. (Individually optimal index-priced contracts.) An individually optimal index-priced contract between shipper i and carrier j on lane ℓ is pegged one-to-one to the spot rate and internalizes shipper's match-specific gain in the following way

$$p_{ij}(\tilde{p}_t) = \begin{cases} b_{ij}^0 + \frac{\text{“incentive”}}{z_{ij}} + \tilde{p}_t & \text{if carrier } j \text{ accepts} \\ b_{ij}^0 & \text{if carrier } j \text{ rejects,} \\ \text{“screening”} \end{cases}$$

where b_{ij}^0 is a fixed fee on which carrier j bids in shipper i 's auction on lane ℓ .

8.2 Welfare comparison

We calculate the market-level welfare of the current institution and the two alternatives on each cluster of lanes for a full-year period in each market phase. Following Acocella, Caplice, and Shef (2020), we take the period from April 15, 2016 to April 14, 2017 as a full year in a tight market, and the period from October 1, 2017 to September 30, 2018 as a full year in a soft market. Our welfare calculation involves two steps. First, we recover the underlying demand and supply factors using the market equilibrium condition in Equation (8). Within each cluster of lanes and each market phase, we recover L relationships, all formed in $t = 0$. For each week t , we recover a direct spot demand D_t and a spot capacity C_t . Second, we calculate the welfare in the current and counterfactual institutions, keeping fixed the underlying demand and supply factors $(L, D_t, C_t)_t$, distribution of match-specific gains and operational costs, while allowing search costs to vary with the equilibrium thickness of the spot market.⁴³

8.2.1 Aggregate welfare

Figure 12 and Table 6 present the per-mile average welfare in three institutions: (i) a centralized spot market for all transactions (“None”), (ii) long-term relationships with fixed-rate contracts coexisting with a spot market (“Fixed-rate”), and (iii) long-term relationships with the individually optimal index-priced contracts coexisting with a spot market (“Index-priced”). To detect the sources of gains and losses, we break down the average welfare into three components: realized match-specific gains, operational costs and search costs. Note that our welfare calculation treats demand for transportation service as inelastic and our welfare measure excludes the benefits to shippers from having their loads transported.⁴⁴ That is, we answer the question on which institution is more efficient at fulfilling a fixed number of loads.

First, we find that centralizing all transactions into a spot platform results in substantial welfare loss from the current institution. The welfare loss is 38 cents/mile in a soft market and 39 cents/mile in a tight market, which are equivalent to 31% and 22% of the median operational cost in the respective market phase. Note that while the reduction in search costs, by 36 cents/mile in a soft market and 21 cents/mile in a tight market, is large, it only benefits those serving in the spot market. Moreover, the reduction in operational costs is small. As a result, the reduction in search and operational costs is far from compensating for the complete loss of match-specific gains from long-term relationships.

Second, individually optimal index-priced contracts can be welfare-improving upon the current fixed-rate contracts, but only in a tight market, when demand for transportation service

⁴³The details of these steps are delegated to Appendix C.2.

⁴⁴That is, we normalize shippers' gains from spot transactions to zero. This is why the average welfare is calculated to be negative.

Figure 12: Welfare comparison across three market institutions

Table 6: Summary of welfare channels

Relationship type	Soft market			Tight market		
	None	Fixed-rate	Index-priced	None	Fixed-rate	Index-priced
Spot rate	1.56	1.63	1.47	2.07	2.06	1.92
Spot share	100%	44%	29%	100%	61%	34%
Search cost	0.36		0.20	0.21		0.26
Match-specific gains		0.57	0.63		0.50	0.64
Operational costs	0.54	0.58	0.63	1.09	1.14	1.23
Search costs	0.01	0.15	0.15	0.12	0.18	0.17
Normalized welfare*	0.53	0.15	0.15	1.21	0.82	0.76

Notes: *Our welfare measure normalizes the shippers' benefits from spot transactions to zero. That is, it includes carriers' costs and match-specific gains and shippers' match-specific gains. All numbers are in \$/mile.

is high. The reason for this difference is that, in a tight market, relationships with fixed-rate contracts face a more serious moral hazard problem. Eliminating moral hazard by individually optimal index-priced contracts thus brings about larger gains in a tight market than it does in a soft market. Specifically, index-priced contracts lead to an increase in match-specific gains relative to the current institution by 11% and 28% in the soft and the tight market respectively. Furthermore, index-priced contracts result in a higher search cost in the spot market, by 20 cents/mile in a soft market and 26 cents/mile in a tight market, but more loads would be accepted within long-term relationships and they entail no search costs. In aggregate, these two forces balance out, resulting in no difference in average search costs between the two types of contracts. There is, however, a more subtle effect of an increase in search costs on aggregate welfare. Higher search costs create a larger wedge in realized operational costs be-

tween carriers servicing in long-term relationships and those servicing in the spot market. This tends to hurt allocative cost efficiency. Across both market phases, operational costs increase by about 9% under index-priced contracts relative to fixed-rate contracts. Another factor contributing to this increase is that carriers in long-term relationships with index-priced contracts fully internalize match-specific gains, thus accepting even when operational costs are high.

8.2.2 Distributional effects

Both counterfactual institutions have large distributional consequences in comparison to the current institution. Figure 13 and Table 7 present the average per-period payoff (\$/mile) of each individual in the market. There are four groups: shippers and carriers who form relationships, and shippers and carriers who only transact in the spot market. We will refer to the latter two groups as spot shippers and spot carriers.

Figure 13: Distributional effects

Table 7: Distributional effects

Relationship type	Soft market			Tight market		
	None	Fixed-rate	Index-priced	None	Fixed-rate	Index-priced
LT relationships						
Shippers		1.02	1.02		1.47	1.46
Carriers		0.71	0.81		0.44	0.65
Spot market						
Shippers	1.56	1.63	1.47	2.07	2.06	1.92
Carriers	0.68	0.50	0.32	0.47	0.36	0.20

Notes: Shippers' benefits from having their loads shipped are normalized to zero. All numbers are in \$/mile.

Overall, institutions with long-term relationships tend to benefit those who manage to form relationships and hurt those who transact only in the spot market. Specifically, better performance of long-term relationships affects the spot market by (i) reducing demand for spot loads and (ii) increasing search costs for spot loads. Both of these channels unambiguously hurt spot carriers. Specifically, relative to a centralized spot platform, the average welfare of a spot carrier reduces by approximately 25% under fixed-rate contracts and 55% under index-priced contracts in both soft and tight market. Spot shippers are affected via equilibrium spot rates, with channel (i) pushing towards lower spot rate and channel (ii) pushing towards higher spot rate. Across the three institutions, the institution with index-priced contracts in long-term relationships has the lowest equilibrium spot rate (\$1.47/ mile), which benefits spot shippers.

Within long-term relationships, index-priced contracts improve the welfare of carriers and not shippers upon fixed-rate contracts. This difference is because the split of surplus from long-term relationships is more favorable to shippers under fixed-rate contracts and more favorable to carriers under index-priced contracts. As the potential surplus from long-term relationships is fully extracted under index-priced contracts, carriers get more information rents in the auctions. Specifically, relative to fixed-rate contracts, index-priced contracts increase the average welfare of carriers in long-term relationships by 14% in a soft market and 48% in a tight market.

8.2.3 Lower-bound comparison to the market-level first-best welfare

In this section, we provide an upper bound on the market-level first-best welfare to benchmark the performance of fixed-rate and index-priced contracts. Since a centralized spot market is unambiguously the worst performing institution, we use it as a baseline for normalization; market-level welfare gain from this baseline will be referred to as market-level surplus.

An upper bound on the market-level first-best welfare. Since search costs are determined endogenously by spot market thickness, it is difficult to calculate market-level first-best welfare. However, an intuitive (strict) upper bound of this welfare measure can be constructed within our framework. First, we fix search costs to the level obtained in a centralized spot market, which is the lowest feasible level of search costs across all market institutions. Then, we exploit index-priced contracts to achieve allocative efficiency, internalizing all match-specific gains from transactions within long-term relationships and the fixed level of search costs for spot transactions.⁴⁵ In other words, our upper bound on the market-level first-best welfare

⁴⁵Index-priced contracts achieve allocative efficiency conditional on a fixed level of search costs by (i) building an aggregate cost curve that internalizes all match-specific gains and search costs and (ii) using a price mechanism to clear the market. Specifically, this aggregate cost curve is made up of the following components. First, each carrier j in a long-term relationship under index-priced contract with shipper i on lane l provides service in period t if its internalized cost is less than the equilibrium spot rate, $c_{j,t} = (c_{ij} + c_{jl}) - \tilde{p}_t$. Second, each spot carrier

Figure 14: Comparison to an upper bound on the market-level first-best welfare

shuts down the negative externalities of transactions in long-term relationships on the spot market. We refer to this upper bound as the unattainable first-best and use it to evaluate the performance of fixed-rate and index-priced contracts. That is, the relative performance of fixed-rate and index-priced contracts to the unattainable first-best provides a lower bound comparison of the welfare of these institutions to the market-level first-best welfare.

Figure 14 plots the ratio of the market-level surplus of fixed-rate and index-priced contracts to the market-level surplus of the unattainable first-best, across our ten clusters of lanes and two market phases. Here, the market-level surplus of an institution is defined as the improvement in market-level welfare from a centralized spot market. Moreover, we order the clusters of lanes by their average distance, since search costs vary less with spot market thickness on longer lanes.

A noteworthy finding is that fixed-rate contracts perform quite well at the market level, capturing around 40% to 70% of the market-level surplus of the unattainable first-best. Index-priced contracts perform similarly to fixed-rate contracts in a soft market and outperform fixed-rate contracts in a tight market. Furthermore, the performance of both fixed-rate and index-priced contracts improves with the distance of lanes, reflecting the lower effect of spot market thickness on search costs on longer lanes. However, the magnitude of this improvement is inflated by the fact that when search costs vary less with spot market thickness, our upper

j^0 provides service in period t if the sum of its operational and search costs is less than the equilibrium spot rate, $c_{j^0,t} + \tau_j \tilde{p}_t$. However, the aggregate cost curve is determined endogenously, shifting upwards due to higher search costs as more match-specific gains in long-term relationships are realized. This is why individually optimal index-priced contracts are not necessarily socially optimal.

bound on the market-level first-best welfare is also tighter.

8.3 Discussion

While our market-level welfare analysis shows that a fully centralized spot market is unambiguously worse than both institutions with long-term relationships, it does not suggest that the role of the spot market should be downgraded. In contrast, the spot market provides an important clearing mechanism for hybrid institutions where long-term relationships and the spot market coexist. In particular, index-priced contracts require reliable measures of spot rates, which will not be available if the spot market is too thin. Moreover, technological advances in the near future can improve the spot market in more ways than by reducing search costs. For example, a digital spot platform can suggest carriers to different lanes to exploit network externalities, thus mitigating the aggregate empty mile problem.

Our comparison of fixed-rate and index-priced contracts demonstrates that the contract design of individual relationships can have market-level consequences. For example, we find that during a soft market and on lanes of shorter distance, fixed-rate contracts, which are suboptimal at the relationship level, generate higher market-level welfare than index-priced contracts. It would be interesting to examine the market-level welfare effects of contracts between these two extremes. In fact, there may be barriers to the implementation of individually optimal index-priced contracts. For example, price uncertainties and the fixed fees that carriers need to pay for rejections would pose concerns for shippers and carriers with budget constraints. Acocella, Caplice, and Sheff (2022a) study the relationship-level effects of index-based contracts that are tuned to practitioners' concerns. Our framework can be used to evaluate the market-level performance of such contracts.

9 Conclusion

This paper studies the interactions and welfare effects of long-term relationships and the spot market. Using detailed data on the US for-hire truckload freight industry, we argue that the two-way crowding-out effects between long-term relationships and the spot market, as hypothesized by Kranton (1996), are present and play a crucial role in this setting. On the one hand, the spot market crowds out long-term relationships by creating a moral hazard problem within relationships. On the other hand, long-term relationships crowd out the spot market by reducing spot market thickness, thereby increasing search costs for spot loads.

Our methodological contribution is a model that captures both the formation of and interactions within long-term relationships. We model relationship formation as an auction and

interactions within the winning relationship as a repeated game. We recover a rich set of model primitives by building on tools from the empirical auction and dynamic discrete choice literature. An empirical challenge of the dynamic discrete choice problem in our setting is that payoff-relevant actions are only partially observed. We tackle this challenge with a novel support-based argument. Specifically, we exploit the sensitivity of the observed decision margin (between “accept” and “reject”) to a running variable (the current spot rate) that affects the unobserved decision margin (between “spot” and “idle”) to pin down the latter decision margin. This identification approach could be generalized to other settings with moral hazard, where actions are naturally not fully observed. Moreover, we develop an identification argument for auctions with two-sided match-specificity. Our argument relies on the observation that the shipper's and the carrier's expected payoffs in a relationship depend on their match-specific gains only through their per-transaction rents. Under empirically plausible conditions, we derive equilibrium conditions that permit a monotone mapping between the carrier's rent and the shipper's rent and pin down this mapping by an approach similar to Guerre, Perrigne, and Vuong (2000). More generally, we demonstrate how in settings where parties share rents, classic insights from the auction literature apply even through there are two latent variables with unrestricted correlation.

Estimating the model, we find that long-term relationships generate large match-specific gains, but realizing more of these gains would come at the cost of a thinner spot market with significantly higher search costs. This market-level tradeoff is a key consideration in evaluating different market institutions. Future technological innovations could either threaten to replace relationships with a more efficient spot market or enhance relationships with more sophisticated contract design. In either case, the welfare effects would be determined by the two-way crowding-out effects that we estimate.

Our counterfactual analysis suggests that the benefits of long-term relationships outweigh their negative externalities. However, this does not mean that the market unambiguously benefits from optimizing the performance of relationships. On the one hand, removing relationships would result in substantial welfare loss, despite achieving the maximal thickness of the spot market. This finding suggests that the dominance of long-term relationships in the current institution is not driven by a coordination failure to form a thick spot market but rather by the large match-specific gains from long-term relationships. On the other hand, optimizing the performance of individual relationships with index-priced contracts leads to only small improvements in market-level welfare and only in periods with high demand. The reason is that such contracts worsen the negative externalities of relationships on the spot market.

While this paper answers key questions about the interactions between long-term relationships and the spot market, there are several possible directions for future research in this

setting. First, one could examine the role of brokers in the current market institution, especially in how shippers and carriers search for and haggle on loads. Second, one could extend the model to allow for interactions across lanes, for example, by letting carriers jointly manage multiple relationships and have access to the spot market on different lanes.

More broadly, our paper is a first step towards understanding the role of technological innovations in shifting the boundary between formal and informal interactions. On the one hand, technological innovations can improve the performance of informal interactions by allowing them to incorporate market information. On the other hand, technological innovations can enable new forms of formal interactions that compete with informal interactions. As technology-driven changes take hold in various industries, understanding the implications of both possibilities is key to anticipating, understanding, and responding to technology's effects.

References

- Acocella, A., Caplice, C., & Shef , Y. (2020). Elephants or gold sh?: An empirical analysis of carrier reciprocity in dynamic freight markets. *Transportation Research Part E: Logistics and Transportation Review*142, 102073.
- Acocella, A., Caplice, C., & Shef , Y. (2022a). The end of set-it-and-forget-it pricing? Index-based freight contract design under uncertainty(MIT Center for Transportation & Logistics Working Paper Series).
- Acocella, A., Caplice, C., & Shef , Y. (2022b). When suppliers ignore their contracts: Quantifying the value of a customer in the buyer-supplier relationship(MIT Center for Transportation & Logistics Working Paper Series).
- Adhvaryu, A., Bassi, V., Nyshadham, A., & Tamayo, J. A. (2020)No line left behind: Assortative matching inside the rm (Working Paper). National Bureau of Economic Research.
- Admati, A. R., & P eiderer, P. (1988). A theory of intraday patterns: Volume and price variability. *The review of nancial studies*1(1), 3–40.
- Allen, J., & Wittwer, M. (2021). Centralizing over-the-counter markets?(Working Paper). Bank of Canada Staff Working Paper.
- Baker, G. P., Gibbons, R., & Murphy, K. J. (1994). Subjective performance measures in optimal incentive contracts. *The Quarterly Journal of Economics*109(4), 1125–1156.
- Baker, G. P., & Hubbard, T. N. (2003). Make versus buy in trucking: Asset ownership, job design, and information. *American Economic Review*93(3), 551–572.
- Baker, G. P., & Hubbard, T. N. (2004). Contractibility and asset ownership: On-board computers and governance in us trucking. *The Quarterly Journal of Economics*119(4), 1443–1479.
- Barron, D., Gibbons, R., Gil, R., & Murphy, K. J. (2020). Relational adaptation under reel authority. *Management Science*66(5), 1868–1889.
- Bester, C. A., & Hansen, C. B. (2016). Grouped effects estimators in xed effects models. *Journal of Econometrics*190(1), 197–208.
- Biais, B., & Green, R. (2019). The microstructure of the bond market in the 20th century. *Review of Economic Dynamics*33, 250–271.
- Bonhomme, S., & Manresa, E. (2015). Grouped patterns of heterogeneity in panel data. *Econometrica*, 83(3), 1147–1184.
- Brancaccio, G., Kalouptsidi, M., Papageorgiou, T., & Rosaia, N. (2020)Search frictions and ef ciency in decentralized transportation markets(Working Paper). National Bureau of Economic Research.
- Brugues, F. (2020). Take the goods and run: Contracting frictions and market power in supply chains(Working Paper). Brown University.

- Buchholz, N. (2022). Spatial equilibrium, search frictions, and dynamic efficiency in the taxi industry. *The Review of Economic Studies* 89(2), 556–591.
- Cajal-Grossi, J., Macchiavello, R., & Noguera, G. (2022) Buyers' sourcing strategies and suppliers' markups in Bangladeshi garments (Working Paper).
- Caliendo, L., Parro, F., Rossi-Hansberg, E., & Sarte, P.-D. (2018). The impact of regional and sectoral productivity changes on the us economy. *The Review of Economic Studies* 85(4), 2042–2096.
- Caplice, C. (2007). Electronic markets for truckload transportation. *Production and Operations Management* 16(4), 423–436.
- Collin-Dufresne, P., Hoffmann, P., & Vogel, S. (2019). Informed traders and dealers in the fx forward market (Working Paper).
- Dugast, J., Üslü, S., & Weill, P.-O. (2019). A theory of participation in otc and centralized markets (NBER Working Paper).
- Elliott, M. (2014). Heterogeneities and the fragility of labor market (Working Paper).
- Frechette, G. R., Lizzeri, A., & Salz, T. (2019). Frictions in a competitive, regulated market: Evidence from taxis. *American Economic Review* 109(8), 2954–92.
- Galenianos, M., & Gavazza, A. (2017). A structural model of the retail market for illicit drugs. *American Economic Review* 107(3), 858–896.
- Gavazza, A. (2010). Asset liquidity and financial contracts: Evidence from aircraft leases. *Journal of Financial Economics* 95(1), 62–84.
- Gavazza, A. (2011). Leasing and secondary markets: Theory and evidence from commercial aircraft. *Journal of Political Economy* 119(2), 325–377.
- Gayle, G.-L., & Miller, R. A. (2015). Identifying and testing models of managerial compensation. *The Review of Economic Studies* 82(3), 1074–1118.
- Gil, R., Kim, M., & Zanarone, G. (2021). Relationships under stress: Relational outsourcing in the US airline industry after the 2008 financial crisis. *Management Science*
- Guerre, E., Perrigne, I., & Vuong, Q. (2000). Optimal nonparametric estimation of first-price auctions. *Econometrica* 68(3), 525–574.
- Harris, A., & Nguyen, T. M. A. (2021). Long-term relationships in the US truckload freight industry (Working Paper).
- Hu, Y., & Xin, Y. (2021). Identification and estimation of dynamic structural models with unobserved choices (Working Paper).
- Hubbard, T. N. (2001). Contractual form and market thickness in trucking. *RAND Journal of Economics* 369–386.
- Kasahara, H., & Shimotsu, K. (2009). Nonparametric identification of finite mixture models of dynamic discrete choices. *Econometrica* 77(1), 135–175.

- Kranton, R. E. (1996). Reciprocal exchange: A Self-Sustaining system. *American Economic Review*, 86(4), 830–851.
- Lagos, R. (2000). An alternative approach to search frictions. *Journal of Political Economy* 108(5), 851–873.
- Lagos, R. (2003). An analysis of the market for taxicab rides in new york city. *International Economic Review* 44(2), 423–434.
- Lee, T., & Wang, C. (2018). Why trade over-the-counter? when investors want price discrimination (Jacobs Levy Equity Management Center for Quantitative Financial Research Paper).
- Macaulay, S. (1963). Non-contractual relations in business: A preliminary study. *The Sociology of Economic Life* (pp. 361–377). Springer.
- Macchiavello, R. (2022). Relational contracts and development. *Annual Review of Economics* 14, 337–362.
- Macchiavello, R., & Morjaria, A. (2015). The value of relationships: Evidence from a supply shock to kenyan rose exports. *American Economic Review* 105(9), 2911–45.
- Macchiavello, R., & Morjaria, A. (2021). Competition and relational contracts in the Rwanda coffee chain. *The Quarterly Journal of Economics* 136(2), 1089–1143.
- Marcus, A. A. (1987). From market dominance to credible commitments: Shippers' strategies in a deregulated trucking environment. *Transportation journal*, 4–15.
- Masten, S. E. (2009). Long-term contracts and short-term commitment: Price determination for heterogeneous freight transactions. *American Law and Economics Review* 1(1), 79–111.
- Ngai, L. R., & Tenreyro, S. (2014). Hot and cold seasons in the housing market. *American Economic Review* 104(12), 3991–4026.
- Nickerson, J. A., & Silverman, B. S. (2003). Why aren't all truck drivers owner-operators? Asset ownership and the employment relation in interstate for-hire trucking. *Journal of Economics & Management Strategy* 12(1), 91–118.
- Ostria, S. J. (2003). Evaluation of US commercial motor carrier industry challenges and opportunities (Technical Report).
- Pagano, M. (1989). Trading volume and asset liquidity. *The Quarterly Journal of Economics* 104(2), 255–274.
- Perrigne, I., & Vuong, Q. (2011). Nonparametric identification of a contract model with adverse selection and moral hazard. *Econometrica* 79(5), 1499–1539.
- Rosaia, N. (2020). Competing platforms and transport equilibrium: Evidence from New York City (Working Paper). Harvard University.
- Rose, N. L. (1985). The incidence of regulatory rents in the motor carrier industry. *The RAND Journal of Economics* 299–318.

- Rose, N. L. (1987). Labor rent sharing and regulation: Evidence from the trucking industry. *Journal of Political Economy* 95(6), 1146–1178.
- Rust, J. (1994). Structural estimation of Markov decision processes. *Handbook of econometrics* 4, 3081–3143.
- Startz, M. (2021). The value of face-to-face: Search and contracting problems in nigerian trade (PEDL Research Papers).
- Train, K. E. (2008). EM algorithms for nonparametric estimation of mixing distributions. *Journal of Choice Modelling* 1(1), 40–69.
- Tunca, T. I., & Zenios, S. A. (2006). Supply auctions and relational contracts for procurement. *Manufacturing & Service Operations Management* 6(1), 43–67.
- Williams, N., & Murray, D. (2020). An analysis of the operational costs of trucking: 2020 update (Technical Report). American Transportation Research Institute.
- Yang, R. (2021). (Don't) take me home: Home bias and the effect of self-driving trucks on interstate trade (Working Paper).
- Ying, J. S. (1990). The inefficiency of regulating a competitive industry: Productivity gains in trucking following reform. *The Review of Economics and Statistics* 72(1), 191–201.
- Zahur, N. B. (2022). Long-term contracts and efficiency in the liquefied natural gas industry (Working Paper).

A Data Construction

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

Contract rates. Shippers seeking long-term relationships define lanes at geographical levels finer than KMA to KMA, sometimes as fine as warehouse-to-warehouse. Shippers can also bundle origin-destination pairs with close proximity as a lane, using the same contract. This means that if a contract specifies a linehaul rate (total payment for a trip) on such a lane, the carrier's per-mile payment would vary with the distance of specific trips. On the other hand, if a contract specifies a per-mile rate, the carrier's total payment would vary with trips' distance. 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.

Figure 15: Number of identified auction and demotion events

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

Auction events. Our data do not include records of auctions. However, we can observe when contract rates change. If we observe at least three changes in contract rates within a date-range from the previous date-range, we assign an auction event to the beginning of the current date-range. The secondary indicator of auction events is when a completely new carrier replaces the previous primary carrier. There is a tradeoff in using this indicator. On the one hand, not using this indicator risks missing some auction events because carriers sometimes reuse their bids. On the other hand, using this indicator risks assigning an auction event to what is actually a demotion event, since some backup carriers may not appear in the routing guide. We perceive the second risk to be smaller and use both indicators to detect auction events.

Demotion events. A demotion event is an instance where the current primary carrier is replaced by a different 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 indicators of auction events.

Figure 15 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. Additionally, there are some 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 show a correlation with identified auction events. This suggests that our data construction does a reasonable job at separating the two types of event.

Carrier's types. There are two public identifying code systems for carriers: the Standard Carrier Alpha Code (SCAC), maintained by the National Motor Freight Traffic Association (NMFTA) and US DOT for carrier registration at the Department of Transportation. We map the SCAC variable in our data set to US DOT codes using a conversion table from the NMFTA. We then map US DOT codes to carriers' registration at the Department of Transportation for the year 2020. This method matches 90% of carriers in our data set to five types: brokers (B), small asset-owners (SC), large asset-owners (LC), broker/small asset-owners (B-SC), broker/large asset-owners (B-LC). The latter two groups are for carriers with multiple divisions.

Figure 16 plots the acceptance probabilities and sensitivities to spot rates by carrier type. It shows that brokers (solid black) are more likely to accept loads, but also more sensitive to spot rates. The first pattern is likely due to brokers' having more flexibility, with a large pool of carriers to draw from in the spot market. The second pattern is likely due to brokers' costs and thus profit margins on contracted loads being directly tied to spot rates. Moreover, we observe that carriers that are both brokers and asset owners behave similarly to asset owners. In this paper, we drop carriers that are identified as brokers (B) from our data set.

B Omitted proofs

B.1 Properties of full compensation and acceptance schedules

Our identification argument relies on the observation that under Assumption 2 and Assumption 3, acceptance schedules have well-defined, distinct “jump” points. This section proves the key properties of the full compensation that give rise to these “jump” points.

When proving the properties of a single relationship, we drop the dependence of notation on the carrier's rent and subscript i, j, ℓ for shipper, carrier and lane for ease of notation. Recall that the full compensation includes the carrier's rent inclusive of savings on search cost and a dynamic incentive component

$$\bar{p}(R_{t-1}, \tilde{p}_t) = \underbrace{\frac{c_t + p_t}{z}}_{\text{carrier's transformed rent}} + \underbrace{\frac{1}{z}}_{\text{dynamic incentive}} [V(R_{t-1}, \tilde{p}_t) - V(R_{t-1} + (1 - \beta)R_{t-1}, \tilde{p}_t)],$$

and transformed cost $\tilde{c}_t = c_t + \frac{c_t + p_t}{z}$ has distribution $\text{Normal}(\tilde{c}^c, \tilde{c}^v)$, denoted by \tilde{F} . The Bellman's

Figure 16: Acceptance tendency across carriers' types

Notes: We run Logit regression $P(d_{ij,t} = \text{accepted}) = \beta_{ij,0} + \beta_{ij,1}(\tilde{p}_t - p_{ij})$ where i denotes the shipper, j denotes the primary carrier, \cdot denote the lane and t denotes the period in their relationship. The left panel plots the quantile distribution of predicted acceptance probabilities across all relationships by carrier type at $\tilde{p}_t = p_{ij}$. The right panel plots the quantile distribution of the decrease in predicted accepted probabilities across all relationships by carrier type when \tilde{p}_t increases by one standard deviation from p_{ij} .

equation of the carrier gives continuation value

$$V(R_{t-1}, \tilde{p}_{t-1}) = E_{\tilde{p}_t}[(1 - \alpha_0(R_{t-1}, \tilde{p}_t))V(\tilde{p}_t) | \tilde{p}_{t-1}] + E_{\tilde{p}_t}[\alpha_0(R_{t-1}, \tilde{p}_t) f(1 - \alpha_0(R_{t-1}, \tilde{p}_t)) h(\tilde{p}(R_{t-1}, \tilde{p}_t), \tilde{p}_t) + V(R_{t-1} + (1 - \alpha_0(R_{t-1}, \tilde{p}_t)), \tilde{p}_t) g | \tilde{p}_{t-1}], \quad (15)$$

where

$$h(\tilde{p}(R_{t-1}, \tilde{p}_t), \tilde{p}_t) = 1f_{\tilde{p}_t < \tilde{p}(R_{t-1}, \tilde{p}_t)} \tilde{p}(R_{t-1}, \tilde{p}_t) g \tilde{F}(\tilde{p}(R_{t-1}, \tilde{p}_t)) (\tilde{p}(R_{t-1}, \tilde{p}_t) - E[\tilde{c}_j \tilde{c}_i | \tilde{p}(R_{t-1}, \tilde{p}_t)]) + 1f_{\tilde{p}_t > \tilde{p}(R_{t-1}, \tilde{p}_t)} \tilde{p}(R_{t-1}, \tilde{p}_t) g \tilde{F}(\tilde{p}_t) (\tilde{p}_t - E[\tilde{c}_j \tilde{c}_i | \tilde{p}_t]) \quad (16)$$

is an expression capturing the carrier's current payoff and the gain in continuation value from an acceptance in the current period.

We use a series of lemmas to show that the full acceptance schedules of relationships with different carriers' rents have well-defined, distinct "jump" points.

Lemma B.1. Suppose that Assumption 3 holds. For every level of carrier's transformed rent $p + \epsilon > 0$ and rejection state R_{t-1} , there exists $\hat{p}_{low}, \hat{p}_{high} \in \mathbb{R}$ such that $\tilde{p}(R_{t-1}, \hat{p}_{low}) = p + \epsilon$ and $\tilde{p}(R_{t-1}, \tilde{p}_t) < p + \epsilon$ for all $\tilde{p}_t > \hat{p}_{high}$.

Proof. Under Assumption 3 (iii), $\bar{p}(R_{t-1}, \hat{p}_{low}) = \hat{p}_{low}$ for all $\hat{p}_{low} \geq p + \tilde{c}_t$. To show the existence of \hat{p}_{high} , it suffices to 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 $p + \tilde{c}_t$, while not being in a relationship only gives the carrier the option to accept a spot load, which gives a payoff of $\tilde{p}_t - \tilde{c}_t$, or to remain idle and get zero. Thus, the continuation value of the carrier at any state is bounded above by the continuation value were the carrier to never be demoted, and bounded below by the continuation value were the carrier to be demoted immediately. It follows that

$$\begin{aligned}
 V(R_{t-1}, A, \tilde{p}_t) &= V(R_{t-1} + (1 - \alpha)\tilde{p}_t, \tilde{p}_t) \\
 &= (1 - \alpha) \mathbb{E}[\max\{p + \tilde{c}_t, \tilde{p}_t - \tilde{c}_t, 0\} | \tilde{p}_t] + \alpha \mathbb{E}[\max\{\tilde{p}_t - \tilde{c}_t, 0\} | \tilde{p}_t] \\
 &= (1 - \alpha) \mathbb{E}[\max\{p + \tilde{c}_t, \tilde{p}_t - \tilde{c}_t, g - \max\{\tilde{p}_t - \tilde{c}_t, g\}\} | \tilde{p}_t] + p + \tilde{c}_t.
 \end{aligned}$$

Thus, the full compensation schedule at (R_{t-1}, \tilde{p}_t) is bounded above, $\bar{p}(R_{t-1}, \tilde{p}_t) \leq \frac{p + \tilde{c}_t}{1 - \alpha}$, completing the proof of the lemma. \square

Definition 4. (Jump points) Fix search cost c . For each level of carrier's rent $p + c$ 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,

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

By Lemma B.1, p is well defined. Moreover, under Assumption 2 on the cost distribution and Assumption 3 on the continuity of the full acceptance schedule in spot rates, acceptance probability is positive in the left neighborhood of $p(R_{t-1}, j + p)$ and zero to the right of this point. This is why we refer to this point as a jump point. Note that we do not rule out the possibility that the full acceptance schedule equals spot rate at multiple points. We focus on the highest such point, since it has the special property that acceptance probability remains zero for any higher level of spot rate. The next lemma shows that these jump points are ordered by the level of carrier rent.

Lemma B.2. (Order of jump points) Suppose that Assumption 2 and Assumption 3 hold. Fix search cost c and cost distribution F , any level of rejection state R_{t-1} , and carriers' rents $p > p^0 + p^0 = 0$. Then,

$$p(R_{t-1}, j + p) > p(R_{t-1}, j^0 + p^0).$$

Proof. Suppose that $p(R_{t-1j} + p) = p(R_{t-1j}^0 + p^0)$, then at $\tilde{p}_t = p(R_{t-1j}^0 + p^0)$,

$$\bar{p}(R_{t-1}, \tilde{p}_t | j^0 + p^0) = \tilde{p}_t - \bar{p}(R_{t-1}, \tilde{p}_t | j + p),$$

where the last equality follows from the definition of \bar{p} and that $p(R_{t-1j} + p) = \tilde{p}_t$. This yields a contradiction because under monotonicity of the full compensation schedule (Assumption 3), we have $\bar{p}(R_{t-1}, \tilde{p}_t | j + p) > \bar{p}(R_{t-1}, \tilde{p}_t | j^0 + p^0)$. \square

B.2 Identification of the distribution of carriers' rents and costs

This section proves the identification of the distribution \tilde{F} of transformed costs and the distribution $[G^{+p}]^{1:N}$ of winning carriers' rents on a given lane. For ease of notation, we drop the dependence of notation on \cdot .

Lemma B.3. (Identification of the distribution of transformed costs) Suppose that Assumption 2 holds. If there exist two distinct jump points $p(R_{t-1j} + p)$ and $p(R_{t-1j}^0 + p^0)$ observed in either two different relationships or two different rejection states of a carrier, then (\tilde{c}, c) are identified.

Proof. Note that each jump point p gives us a point $(p, \tilde{F}(p))$ on the distribution \tilde{F} , where $\tilde{F}(p)$ is the observed acceptance probability at this jump point. Thus, two distinct jump points give a system of linear equations

$$\begin{aligned} \frac{p(R_{t-1j} + p)}{c} - \tilde{c} &= \frac{1}{\tilde{c}} (\Pr(d_t = \text{accept} | R_{t-1}, \tilde{p}_t = p(R_{t-1j} + p); + p)) \\ \frac{p(R_{t-1j}^0 + p^0)}{c} - \tilde{c} &= \frac{1}{\tilde{c}} \Pr(d_t = \text{accept} | R_{t-1}, \tilde{p}_t = p(R_{t-1j}^0 + p^0); ^0 + p^0), \end{aligned}$$

the right hand side of which are observed. This system pins down (\tilde{c}, c) . \square

Lemma B.4. (Identification of carriers' rents in long relationships) Suppose that Assumption 2 and conditions (ii) and (iii) of Assumption 3 hold. In addition, suppose that search cost and the distribution F of operational costs are identified. If a relationship has duration $T \geq 1$, the rent level $+ p$ of the carrier in this relationship is identified.

Proof. Under Assumption 2 and that $T \geq 1$, the acceptance schedule is fully observed. Note that the identification of the carrier rent immediately follows from the monotonicity of the full compensation schedule in carrier rent (Assumption 3 (ii)). Here we present a direct proof that does not rely on monotonicity.

That the acceptance schedule is fully observed means that at any state (R_{t-1}, \tilde{p}_t) in which

$\bar{p}(R_{t-1}, \tilde{p}_t) = \tilde{p}_t$, the value of the full compensation at that state is identified by

$$\bar{p}(R_{t-1}, \tilde{p}_t) = \tilde{F}^{-1}(\Pr(d_t = \text{accept} | R_{t-1}, \tilde{p}_t; \gamma + p)).$$

It follows that at any state (R_{t-1}, \tilde{p}_t) , $h(\bar{p}(R_{t-1}, \tilde{p}_t), \tilde{p}_t)$ in Equation (16) is identified. Thus, we can define a mapping $h : V \rightarrow V$, where each element $(V(R_{t-1}, \tilde{p}_t))_{R_{t-1}, \tilde{p}_t}$ of V satisfies that at each state (R_{t-1}, \tilde{p}_t) ,

$$E[V(\tilde{p}_t) | \tilde{p}_{t-1}] = V(R_{t-1}, \tilde{p}_{t-1}) + \frac{1}{\gamma} p(R_{t-1}).$$

Under Assumption 3, this means that V is bounded and that it contains the solution to the Bellman equation. It is straightforward to show that h is a contraction mapping by verifying that it satisfies Blackwell's sufficient conditions. Thus, it has a unique fixed point, which is also the collection of continuation values in this relationship. We can find this fixed point by

$$(V(R_{t-1}, \tilde{p}_{t-1}))_{R_{t-1}, \tilde{p}_{t-1}} = \lim_{k \rightarrow \infty} h^k(E[V(\tilde{p}_t) | \tilde{p}_{t-1}])_{R_{t-1}, \tilde{p}_{t-1}}.$$

That is, Equation (15), the observed acceptance schedule and \tilde{F} pin down V . Finally, it holds at the jump point that

$$p(R_{t-1}) = \gamma + p + \frac{1}{\gamma} [V(R_{t-1}, p(R_{t-1})) - V(R_{t-1} + (1 - \gamma), p(R_{t-1}))].$$

This pins down carrier rent $\gamma + p$. □

While the identification proof for long relationships helps demonstrate the source of identification power in our setting, we need to develop an argument for the identification of the mixture of carrier rent that includes relationships of all lengths. This argument relies on jump points being strictly monotone in carrier rent. We reproduce the statement and provide the proof below.

Lemma 4. (Identification of the distribution of carrier rent) Suppose that Assumption 2 and Assumption 3 hold. In addition, suppose that the shipper's incentive scheme search cost γ , and the distribution F of operational costs are identified. If the distribution $[G^{+p}]^{1:N}$ of the rents of winning carriers permits an absolutely continuous density, then it is nonparametrically identified.

Proof of Lemma 4. Fix a rejection state R_{t-1} . We exploit the following equality

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

where the joint distribution of carriers' acceptance, rejection states and spot rates, both unconditional and conditional on carriers' rents, are either directly observed or identified. Our task is to identify the mixture $[G^{+p}]^{1:N}$.

Another key property is that beyond the "jump" points, acceptance probability equals zero,

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

Take any two distributions $[G^{+p}]^{1:N}$ and $[\hat{G}^{+p}]^{1:N}$ with absolutely continuous densities $[g^{+p}]^{1:N}$ and $[\hat{g}^{+p}]^{1:N}$ that are not everywhere the same. Let $\bar{r} = \inf\{r^0 : [g^{+p}]^{1:N}(r) = [g^{+p}]^{1:N}(r^0), \forall r > r^0\}$ and suppose, without loss of generality, that $[g^{+p}]^{1:N}(\bar{r}) > [\hat{g}^{+p}]^{1:N}(\bar{r})$. The continuity of $[g^{+p}]^{1:N}$ and $[\hat{g}^{+p}]^{1:N}$ further implies that for some $\epsilon > 0$, $[g^{+p}]^{1:N}(\bar{r}) > [\hat{g}^{+p}]^{1:N}(\bar{r})$ for all $r \in [\bar{r} - \epsilon, \bar{r}]$. Then, it follows from Lemma B.2 that

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

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

B.3 Existence of a symmetric monotone equilibrium

This section constructs a symmetric monotone equilibrium in two steps. First, we construct a monotone equilibrium in a pseudo-game in which only match quality matters. Second, we derive a symmetric monotone equilibrium in the original game from the monotone equilibrium of the pseudo-game.

Assumption 4. There exists $\underline{b} \leq R$ such that for all $b \in \underline{b}, U(R_0, \tilde{p}_0 | r, b) \geq U(R_0, \tilde{p}_0)$ for all $r \geq 0$ and $U(R_0, \tilde{p}_0 | r, b)$ is increasing in $r \geq 0$ and $b \in \underline{b}$.

The intuition for this assumption is that if the shipper's rent is sufficiently high, then raising her rent, the shipper benefits from the carrier having higher rent and thus accepting more frequently. While intuitive, this statement relies on the specifics of how the carrier's rent affects its' path of play and how such path of play is correlated with the realized path of spot rates. The right panel of Figure 22 demonstrates that this assumption is satisfied under our estimated spot process and incentive scheme.

Proposition B.1. Under Assumptions 2, 3 and 4, there exists a symmetric monotone equilibrium.

Proof. We construct a symmetric monotone equilibrium in two steps.

Step 1: A monotone equilibrium of a pseudo-game.

Consider a bidding game where each carrier j has private information about their 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 $b : \theta_{ij} \rightarrow b_{ij}$ such that

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

Note that a relationship strictly benefits the carrier if and only if the carrier's rent is strictly positive. Thus, in this equilibrium, the lowest match quality of a winning carrier gives zero rent to that carrier, $b(\underline{\theta}) = \underline{b}$. 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 it would otherwise strictly benefit from deviating to a lower bid. Denote by $r : \theta_{ij} \rightarrow r_{ij} = b(\theta_{ij})$ the function that maps the carrier's match quality to its rent. We have $r(\underline{\theta}) = 0$, and for all θ , $r(\theta) > 0$ and $b^0(\theta) + r^0(\theta) = 1$. We want to show that r is strictly increasing.

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

$$(N-1) \frac{g(\theta)}{G(\theta)} = \frac{\frac{\partial}{\partial r} V(R_0, \tilde{p}_0 | r = r(\theta))}{V(R_0, \tilde{p}_0 | r = r(\theta)) - E[V(\tilde{p}_1) | \tilde{p}_0]} b^0(\theta)$$

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

$$\begin{aligned} \frac{\frac{\partial}{\partial r} V(R_0, \tilde{p}_0 | r = r(\theta_1))}{V(R_0, \tilde{p}_0 | r = r(\theta_1)) - E[V(\tilde{p}_1) | \tilde{p}_0]} &= \frac{g(\theta_1)}{G(\theta_1)} [b^0(\theta_1)]^{-1} \\ &> \frac{g(\theta_2)}{G(\theta_2)} [b^0(\theta_2)]^{-1} = \frac{\frac{\partial}{\partial r} V(R_0, \tilde{p}_0 | r = r(\theta_2))}{V(R_0, \tilde{p}_0 | r = r(\theta_2)) - E[V(\tilde{p}_1) | \tilde{p}_0]}. \end{aligned}$$

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

Step 2: Symmetric monotone equilibrium.

We now map the monotone equilibrium of the pseudo-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 b , then carrier j has no incentive to deviate from bidding according to b . It remains to show the shipper's selection rule in the pseudo-game is optimal in the original bidding game.

Under Assumption 4 and by the choice of \underline{b} in the pseudo-game, we have for all \underline{p}_0 ,

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

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

C Estimation details

C.1 Estimation of model primitives

Figure 17 presents a roadmap of our estimation procedure.

C.1.1 Estimate the number of bidders

Assume that the number of bidders in an auction is stochastic, $N_a \sim \text{Binomial}(N, q)$. Then the number of effective 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 $\tilde{q} = q(1 - G^{+P}(r))$.

The empirical challenge in estimating (N, \tilde{q}) is that we only observe the number of carriers that receive at least an offer within the auction period. This number, denoted by \hat{n}_a , could be smaller than the number of effective bidders n_a , since low-rank carriers may never receive an offer. We tackle this issue in two steps. First, we take the maximum of \hat{n}_a in all auctions as an estimate of N . Second, we estimate \tilde{q} through a calibration exercise that captures the bias in the number of observed carriers. This exercise simulates a distribution of \hat{n}_a from (N, \tilde{q}) , the total number of offers within each auction, and the estimated probability that a load is rejected conditional on previous rejections and the current spot rate. Matching the mean and variance of this simulated distribution to its empirical counterpart pins down \tilde{q} .

C.1.2 The likelihood contribution of each relationship

For each relationship, we observe the duration of the relationship T_{ij} , and for each period, whether the carrier accepts, the rejection index at the beginning of the period $R_{j,t-1}$, and

Figure 17: Estimation roadmap

Step 1. Estimate instrumental objects

- Calibrate the discount factor β .
- Calibrate the number of effective bidders to be distributed as $\text{Binomial}(N, \tilde{q})$.
- Estimate the spot process by OLS and shippers' strategies by MLE.

Step 2. Estimate cost parameters

Estimate $\beta^c, \tilde{c}_{ia}^c, (p_{ij} + p_{ij} + \dots)_{ij}, T_{ij}, 50$ by MLE.

Outer: $\beta^c \in \{0.1, 0.2, \dots, 2.0\}$.

Inner: For each shipper-lane-auction, estimate the transformed cost shifter \tilde{c}_{ia}^c and the transformed rents of all primary carriers, $(p_{ij} + p_{ij} + \dots)_{ij}$, including both the auction winner and promoted carriers.

Estimate the scale efficiency parameter γ_1 by 2SLS in

$$\tilde{c}_{ia}^c = \frac{\gamma_1 \ln(\text{Volume}^{\text{spot}})}{\text{Distance}} + \gamma_2 1_{\text{tight}}^a + h_3(\text{Distance}) + \beta^c + \tilde{c}_{ia}^c, \quad (17)$$

using the predicted trade flows (Caliendo, Parro, Rossi-Hansberg, & Sarte, 2018) to instrument for $\text{Volume}^{\text{spot}}$. Here h_3 is a polynomial of degree 3.

Extrapolate transformed costs \tilde{c}_{ia}^c from relationships with $T_{ij} = 50$ to all relationships,

$$b_{ia}^c = h_2^0(\text{Rate}^{\text{spot}}, \text{Volume}^{\text{spot}}, \text{Distance}) + \gamma_2 1_{\text{tight}}^a + h_3^0(\text{Distance}).$$

Here h_2^0 and h_3^0 are polynomials of degree 2 and degree 3 respectively.

Decompose each b_{ia}^c into a cost shifter b_{ia}^c and a search cost b_{ij} by normalizing the median operational cost on a lane to \$1.22/mile (\$1.55/mile net of \$0.33/mile fuel surcharge).

Step 3. Estimate the distribution of rents and match-specific gains

Cluster all relationships:

Outer: 10 K-means clusters based on lane characteristics $(\text{Distance}, \text{Rate}^{\text{spot}}, \text{Volume}^{\text{spot}})$.

Inner: 2 market phases (soft, tight) based on the start date of each relationship.

For each sub-cluster:

(i) Estimate the distribution of winning carriers' rents $[G^{+p}]^{1:N}$ by an EM-algorithm.

(ii) Fix the set of median characteristics. Estimate the distribution of shippers' rents from the first-order conditions of carriers' optimal bidding at the percentiles of $[\hat{G}^{+p}]^{1:N}$. Shippers' IR constraint is evaluated at the f th percentile of $[\hat{G}^{+p}]^{1:N}$.

(iii) Estimate the parameters of the fundamental distribution of match-specific gains, G^* , by matching the moments of the distribution of carriers' rents conditional on contract rates.

the mean spot rate in that period \tilde{p}_t . We also observe the standard deviation of spot rates, σ_t . The likelihood contribution of this relationship depends on the parameters of the carrier's transformed costs (\tilde{c}_{ia}^c, c^c) and transformed rent $\tilde{r}_{ij} + p_{ij} + \tilde{r}$ as follows

$$\ln L = \sum_{t=1}^T \sum_{j \in \mathcal{J}} \Pr(d_{j,t} = \text{accept} | R_{j,t-1}, \tilde{p}_t; \tilde{c}_{ia}^c, c^c, \tilde{r}_{ij} + p_{ij} + \tilde{r})^{1-d_{j,t}} \Pr(d_{j,t} = \text{idle} | R_{j,t-1}, \tilde{p}_t; \tilde{c}_{ia}^c, c^c, \tilde{r}_{ij} + p_{ij} + \tilde{r})^{d_{j,t}}$$

where for each t ,

$$\Pr(d_{j,t} = \text{accept} | R_{j,t-1}, \tilde{p}_t; \tilde{c}_{ia}^c, c^c, \tilde{r}_{ij} + p_{ij} + \tilde{r}) = \frac{\bar{p}(R_{j,t-1}, \tilde{p}_t | \tilde{c}_{ia}^c, c^c, \tilde{r}_{ij} + p_{ij} + \tilde{r})}{\bar{p}(R_{j,t-1}, \tilde{p}_t | \tilde{c}_{ia}^c, c^c, \tilde{r}_{ij} + p_{ij} + \tilde{r}) + \bar{p}(R_{j,t-1}, \tilde{p}_t | \tilde{c}_{ia}^c, c^c, \tilde{r}_{ij} + p_{ij} + \tilde{r})} \quad (18)$$

$$= \frac{\bar{p}(R_{j,t-1}, \tilde{p}_t | \tilde{c}_{ia}^c, c^c, \tilde{r}_{ij} + p_{ij} + \tilde{r})}{\bar{p}(R_{j,t-1}, \tilde{p}_t | \tilde{c}_{ia}^c, c^c, \tilde{r}_{ij} + p_{ij} + \tilde{r}) + \bar{p}(R_{j,t-1}, \tilde{p}_t | \tilde{c}_{ia}^c, c^c, \tilde{r}_{ij} + p_{ij} + \tilde{r})} \quad (19)$$

Given a set of parameter values ($\tilde{c}_{ia}^c, c^c, \tilde{r}_{ij} + p_{ij} + \tilde{r}$), the optimal strategy and value function of carrier j are obtained through an iterative procedure, using three conditions from the carrier's dynamic programming problem. For ease of notation, we will drop the dependence of these conditions on i, j, \tilde{r} and the parameter values.

1) Optimality condition

$$\bar{p}(R_{t-1}, \tilde{p}_t) = \tilde{r} + p + \frac{1}{1 + \rho} (V(R_{t-1}, A, \tilde{p}_t) - V(R_{t-1} + (1 + \rho), \tilde{p}_t)).$$

2) Value function

$$V(R_{t-1}, \tilde{p}_{t-1}) = E_{\tilde{p}_t} [(1 + \rho)(R_{t-1}, \tilde{p}_t) \underline{V}(\tilde{p}_t) | \tilde{p}_{t-1}] + E_{\tilde{p}_t} [R_{t-1}, \tilde{p}_t] (1 + \rho) h(\bar{p}(R_{t-1}, \tilde{p}_t), \tilde{p}_t + \tilde{r}) + V(R_{t-1} + (1 + \rho), \tilde{p}_t) g | \tilde{p}_{t-1}].$$

3) Carrier's expected payoff from the outside option

$$\underline{V}(\tilde{p}_t) = (1 + \rho) E_t [\tilde{F}(\tilde{p}_t + \tilde{r}) (\tilde{p}_t + \tilde{r} - E[\tilde{c}_t | \tilde{c}_t - \tilde{p}_t + \tilde{r}])] + E_{\tilde{p}_{t+1}} [\underline{V}(\tilde{p}_{t+1}) | \tilde{p}_t].$$

Our procedure initializes the carrier value function by its expected payoff from the spot market. In each iteration, we update the full compensation \bar{p} by condition 1) and the value function V by condition 2). The procedure ends when V converges.

Selection on survivals Focusing on long-lasting relationships has two potential selection issues: relationships that last long tend to have either (i) high match-specific gains or (ii) low cost draws. Since the carrier's match-specific gain is a free parameter in an individual relationship when we estimate cost parameters, selection due to (i) is not a concern. Selection due to (ii) can be corrected by conditioning on "surviving" in our construction of the likelihood function. Specifically, acceptance probability at each state will be replaced by

$$\Pr(d_{j,t} = \text{accept} | R_{j,t-1}, \tilde{p}_t, \text{surviving}) = \frac{\phi(R_{j,t-1}, \tilde{p}_t) \Pr(d_{j,t} = \text{accept} | R_{j,t-1}, \tilde{p}_t)}{\phi(R_{j,t-1}, \tilde{p}_t) \Pr(d_{j,t} = \text{accept} | R_{j,t-1}, \tilde{p}_t) + \phi(R_{j,t-1} + (1 - \alpha), \tilde{p}_t) \Pr(d_{j,t} = \text{reject} | R_{j,t-1}, \tilde{p}_t)}$$

if $t < T_{ij}$, and

$$\Pr(d_{j,t} = \text{accept} | R_{j,t-1}, \tilde{p}_t, \text{non-surviving}) = \frac{[1 - \phi(R_{j,t-1}, \tilde{p}_t)] \Pr(d_{j,t} = \text{accept} | R_{j,t-1}, \tilde{p}_t)}{[1 - \phi(R_{j,t-1}, \tilde{p}_t)] \Pr(d_{j,t} = \text{accept} | R_{j,t-1}, \tilde{p}_t) + [1 - \phi(R_{j,t-1} + (1 - \alpha), \tilde{p}_t)] \Pr(d_{j,t} = \text{reject} | R_{j,t-1}, \tilde{p}_t)} \quad \text{CE}$$

if $t = T_{ij}$ and the relationship is ended because of a demotion. We will incorporate these corrections into the next iteration of our estimation. However, since the contribution of an acceptance or rejection to the probability of demotion is small, we expect negligible changes in the results due to these corrections.

C.1.3 The EM-algorithm

We approximate the (continuous) distribution of carriers' rents by a mixture of $K = 5$ Normal distributions. Thus, the parameters to estimate are the mean and variance of each distribution, $(\mu_k, \sigma_k)_{k=1}^K$, and their shares, $(\pi_k)_{k=1}^K$. For estimation, we adapt an EM algorithm by Train (2008): the M-step integrates the likelihood function over these Normal distributions, and the E-step updates their means, variances, and shares. To speed up the integration step, we discretize carrier rent into a grid, $f \in \{0.0, 0.1, \dots, 5.0\}$, and perform linear interpolation on these grid points. In other words, the likelihood contribution of each relationship at each grid point is calculated only once, and when the distribution of carrier rent is updated, we only need to update the weights being put on these grid points.

Treatment of heterogeneity within sub-clusters. In estimating the distribution of carrier rent in each sub-cluster, we keep observable characteristics relationship-specific rather than

using a representative set of characteristics for all relationships in the sub-cluster. The reason for our decision is to avoid in ating the heterogeneity of the estimated distribution of carrier rent. Via an auction approach, such in ated heterogeneity would result in an upward bias of the estimated distribution of shipper rent. Representative (median) characteristics are only used in subsequent steps, where we need to estimate the expected payoffs of shippers and carriers conditional on their rents at the auction stage.

C.1.4 Carriers' bidding function

Our empirical model allows for two sources of randomness in the number of effective bidders: (i) the number of carriers who submit a bid, $N_a \sim \text{Binomial}(N, q)$, and (ii) the number of carriers who pass the shippers' individual rationality constraint, $n_a \sim \text{Binomial}(N, \tilde{q})$. We will show that a carrier's bidding problem can be rewritten as follows,

$$\max_r [\tilde{G}^{+p}(r)]^{N-1} (V(R_0, \tilde{p}_0) - b_r(r)) E[V(\tilde{p}_1) | \tilde{p}_0], \quad (20)$$

where carriers' probability of winning, $[\tilde{G}^{+p}(r)]^{N-1}$, can be estimated from the distribution of winning carriers' rents.

Recall that $[G^{+p}(r)]^n$ is the distribution of winning carriers' rents in auctions with n effective bidders. The distribution of winning carriers' rents estimated in C.1.3 aggregates auctions with different numbers of bidders, conditional on there being at least one that passes shippers' individual rationality constraint:

$$\frac{\sum_{n=1}^N \binom{N}{n} \tilde{q}^n (1 - \tilde{q})^{N-n} \frac{G^{+p}(r) - G^{+p}(\underline{r})}{1 - G^{+p}(\underline{r})}^{n-1}}{\sum_{n=1}^N \binom{N}{n} \tilde{q}^n (1 - \tilde{q})^{N-n}} = \frac{\int_0^{\tilde{G}^{+p}(r)} \frac{G^{+p}(r) - G^{+p}(\underline{r})}{1 - G^{+p}(\underline{r})}^{n-1} (1 - \tilde{q})^N}{1 - (1 - \tilde{q})^N}.$$

Given \tilde{q} from Appendix C.1.1, the above equation pins down $\tilde{G}^{+p}(r) = 1 - \tilde{q} + \tilde{q} \frac{G^{+p}(r) - G^{+p}(\underline{r})}{1 - G^{+p}(\underline{r})}$. Moreover, this gives carriers' probability of winning conditional on their rent, since

$$\sum_{n=1}^N \binom{N}{n} \tilde{q}^{n-1} (1 - \tilde{q})^{N-n} \frac{G^{+p}(r) - G^{+p}(\underline{r})}{1 - G^{+p}(\underline{r})}^{n-1} = [\tilde{G}^{+p}(r)]^{N-1}.$$

To estimate carriers' bidding function b_r , we estimate b_r^0 from the first-order-condition of (20) and the initial condition $b_r(\underline{r})$ from the binding individual rationality constraint of shippers. These first-order-conditions are evaluated at the percentiles of the distribution of winning

carriers' rents estimated in C.1.3, and the lowest rent-type τ is the τ th percentile of this distribution.

C.1.5 Estimate the fundamental distribution of match-specific gains

We parameterize the underlying distribution of match-specific gains G_{ij} by

$$G_{ij} \sim \text{Normal}(\mu, \sigma^2)$$

Given the lowest type τ and bidding function b_r estimated in the previous steps, we can simulate from G_{ij} the joint distribution of match-specific gains, contract rates, rents and bids of observed relationships. We estimate the parameters of G_{ij} by matching simulated moments to the empirical moments of observed relationships.

Specifically, consider two bins of contract rates divided by the median contract rate p_{med} . We use the first and second moments of the distribution of carrier rent in these two bins, $[G_{ij} | p \leq p_{med}]^{1:N}$ and $[G_{ij} | p > p_{med}]^{1:N}$, and the first and second moments of the distribution of contract rates, $[G^p]^{1:N}$. For the empirical moments, we use the empirical distribution of contract rates and use an EM-algorithm to obtain the distributions of carrier rent within each bin of contract rates.

C.2 Details of counterfactual analysis

C.2.1 Estimate cluster-specific market shocks

For each lane-specific cluster and market phase, we combine our sample of long-term relationships and DAT data on the spot market to construct market-level demand and supply factors in a full year. Our sample of long-term relationships gives the total number of loads demanded by the shippers (\hat{L}_t), those loads that are accepted by the primary carrier ($\hat{L}_t^{primary}$), those that are accepted by a backup carrier (\hat{L}_t^{backup}), and those that are rejected by all carriers and fulfilled in the spot market. Our spot market data gives the spot rate \tilde{p}_t and the number of spot loads \hat{S}_t . We scale the number of loads in our sample on long-term relationships to reflect the aggregate market share of long-term relationships, obtaining a series $(\hat{L}_t, \hat{L}_t^{primary}, \hat{L}_t^{backup}, \hat{S}_t, \tilde{p}_t)_{t=1}^{52}$.

From the series of observed loads and spot rates, we recover a series of demand for long-term relationships, direct spot demand and spot capacity, $(L_t, D_t, C_t)_{t=1}^{52}$ that are consistent with the market equilibrium condition in Equation (8). There are two empirical issues: first, our sample can provide noisy estimates of weekly volumes within long-term relationships and second, our model abstracts from backup carriers. We address these issues by assuming that

the demand for long-term relationships is constant within a market phase and counting loads accepted by backup carriers as spot loads. Specifically, we set $L_t = L = \frac{1}{52} \sum_{t=1}^{52} \hat{L}_t$ and assume that all relationships start at $t = 0$. This allows us to capture some correlation in the evolution of the relationship status and the market condition. We then use our estimates of the incentive scheme, distribution of match-specific gains, search and operational costs to estimate the number of loads accepted by the primary carriers (L_t^{primary}) and the number of carriers that reject loads offered within relationships to service the spot market (L_t^{spot}). We also estimate the number of loads accepted by backup carriers (L_t^{backup}) and count these loads towards spot loads.

These volume estimates, the estimated model primitives, and the observed spot rate allow us to pin down direct spot demand and spot capacity through the market equilibrium condition

$$L + D_t = L \frac{Z}{F(\bar{p}_t)} \left\{ \frac{L_t^{\text{primary}}}{L_t^{\text{primary}}} \right\} + \left[C_t F(\bar{p}_t) + \frac{L_t^{\text{spot}}}{L_t^{\text{backup}}} \frac{Z}{F(\bar{p}_t)} \right] \left\{ \frac{L_t^{\text{spot}}}{L_t^{\text{backup}}} \right\}.$$

$$S_t = \hat{S}_t + L_t^{\text{backup}}$$

Specifically,

$$D_t = L_t^{\text{primary}} + L_t^{\text{backup}} + \hat{S}_t L \quad \text{and} \quad C_t = \frac{\hat{S}_t + L_t^{\text{backup}}}{F(\bar{p}_t)} \frac{L_t^{\text{spot}}}{L_t^{\text{spot}}}.$$

In our counterfactual analysis, we keep fixed $(L, D_t, C_t)_{t=1}^{52}$.

C.2.2 Market-level welfare of a centralized spot market

In a centralized spot market, demand for long-term relationships is combined with direct spot demand, and all carriers in the market make up spot capacity. Thus, the equilibrium spot rate in each period is pinned down by

$$L + D_t = \left(\frac{L + C_t}{S_t^{\text{centralized}}} F(\bar{p}_t) \right)^{-1},$$

where

$$S_t^{\text{centralized}} = \frac{P_t^{\text{centralized}}}{\alpha_1 \ln \left(\frac{P_t^{\text{centralized}}}{S_t} \right) + \text{Distance}}$$

is the equilibrium search cost in a centralized spot market. Notice how the new level of search cost depends only on the scale efficiency parameter α_1 , distance and how much spot market volume has scaled up due to the centralization of all transactions into the spot market. Normal-

izing shippers' gains from having their loads shipped to zero, we obtain the following measure of aggregate welfare of a centralized spot market in period t ,

$$W_t^1 = \int_0^X (L + D_t) (c_t - E[c_t] + \tilde{p}_t) d[G]^{1:N}(\cdot).$$

C.2.3 Market-level welfare of index-priced contracts

Denote by c_{ij}^2 the equilibrium search cost under index-priced contracts. In addition, recall that $\tilde{q}_{ij} = c_{ij} + c_{ij}^2$ denotes the match-quality of carrier i and shipper j , excluding savings on search costs. Under index-priced contracts, any relationship with $\tilde{q}_{ij} + c_{ij}^2 > 0$ generates surplus over spot transactions. That is, the lowest match-quality in a relationship is $\tilde{q}_{ij} = c_{ij}^2$. Moreover, carriers in long-term relationships with individually optimal index-priced contracts might reject when costs are high, but never reject in order to service the spot market. Specifically, carrier j rejects if and only if $\tilde{q}_{ij} + \tilde{p}_t < c_t$. Thus, the equilibrium spot rate in each period is pinned down by

$$L + D_t = L \int_0^Z F(\tilde{q}_{ij} + \tilde{p}_t) d[G]^{1:N}(\cdot) + C_t F(\tilde{p}_t) \frac{S_{index}}{S_t},$$

where

$$c_{ij}^2 = \frac{P \ln \frac{S_{index}}{S_t}}{\text{Distance}}.$$

The aggregate welfare under index-priced contracts in period t is

$$W_t^2 = L \int_0^Z F(\tilde{q}_{ij} + \tilde{p}_t) (E[c_t] + \tilde{p}_t) d[G]^{1:N}(\cdot) + C_t F(\tilde{p}_t) (c_t - E[c_t] + \tilde{p}_t).$$

To estimate the split of relationship surplus between the shipper and the carrier, we rely on two observations. First, under an individually optimal index-priced contract, the full surplus from an individual relationship is realized. Second, the shipper's surplus is precisely the fixed fee b_{ij}^0 that the carrier bids on. Denote by $\text{Surplus}(\tilde{p}_0, \tilde{q}_{ij})$ the total expected surplus of the relationship between shipper i and carrier j with match quality \tilde{q}_{ij} and initial spot rate \tilde{p}_0 . There exists a symmetric monotone bidding equation $b^0 : \text{Surplus}(\tilde{p}_0, \tilde{q}_{ij}) \rightarrow b_{ij}^0$ mapping each level of expected surplus to a bid. Specifically, b^0 satisfies that for all $\tilde{q}_{ij} > 0$,

$$b^0(\tilde{q}_{ij}) = \text{Surplus}(\tilde{p}_0, \tilde{q}_{ij}) \frac{R \int_0^Z \text{Surplus}(\tilde{p}_0, \tilde{q}_{ij}) d[G]^{1:N}(\cdot)}{[G]^{1:N}(\tilde{q}_{ij})}.$$

This in turn pins down the expected surplus of the shipper and the carrier in each relationship.

C.2.4 An upper bound on market-level first-best welfare

Our upper bound on the market-level first-best welfare in a period uses the same formula as the welfare under index-priced contracts but replaces the search cost under index-priced contracts with the search cost of a centralized spot market,

$$\bar{W}_t = L \int_0^Z F(\tilde{p}_t + \tilde{p}_t) (E[c_t | c_t + \tilde{p}_t]) d[G]^{1:N}(\cdot) + C_t F(\tilde{p}_t - 1) (E[c_t | c_t - \tilde{p}_t - 1]).$$

In each counterfactual, we report the welfare averaged over time: $W^0 = \frac{1}{52} \sum_t W_t^0$ for fixed-rate contracts (baseline), $W^1 = \frac{1}{52} \sum_t W_t^1$ for a centralized spot market, $W^2 = \frac{1}{52} \sum_t W_t^2$ for index-priced contracts and $\bar{W} = \frac{1}{52} \sum_t \bar{W}_t$ for the unattainable market-level first-best welfare. The performance of fixed-rate and index-priced contracts are measured as the ratio of their surplus relative to a centralized spot market over the surplus of the unattainable first-best relative to a centralized spot market, $\frac{W^0 - W^1}{\bar{W} - W^1}$ and $\frac{W^2 - W^1}{\bar{W} - W^1}$ respectively.

C.3 Lane-specific shares of spot volumes

We construct fitted values from the coefficient estimate $\hat{\alpha}_1 = 0.345$ in the (IV) specification in Table 4 and the fact that the aggregate share of spot volume is 20%. Denote by SpotShare the share of spot volume on lane ℓ . We have

$$\frac{\text{Volume}^{\text{spot}}}{\text{Volume}^{\text{LT}}} = \exp(\tilde{\alpha}_0) \frac{\text{Volume}^{\text{spot}} + \text{Volume}^{\text{LT}}}{\text{Volume}^{\text{LT}}}^{0.345} \quad \text{and} \quad \sum_{\ell} \frac{\text{Volume}^{\text{spot}}_{\ell}}{(\text{Volume}^{\text{LT}}_{\ell} + \text{Volume}^{\text{spot}}_{\ell})} = 20\%,$$

where $\text{Volume}^{\text{LT}}$ is the population long-term relationship volume on lane ℓ , and $\tilde{\alpha}_0$ is the scaling parameter that we need to calibrate. Rewriting the above equations gives,

$$\frac{\text{SpotShare}}{1 - \text{SpotShare}} = \exp(\tilde{\alpha}_0) \frac{\text{Volume}^{\text{spot}}}{\text{SpotShare}}^{0.345} \quad \text{and} \quad \sum_{\ell} \frac{\text{Volume}^{\text{spot}}_{\ell}}{\text{Volume}^{\text{spot}}_{\ell} + \text{SpotShare}_{\ell}} = 20\%,$$

which we use to calibrate $\tilde{\alpha}_0$ and $(\text{SpotShare})_{\ell}$ from the observed spot volume $(\text{Volume}^{\text{spot}})_{\ell}$.

Table 8: Estimates of the relational incentive scheme

Rejection index I_R			Demotion probability $1 - \rho$		
Parameter	Estimate	95% CI	Variable	Estimate	95% CI
R_0	0.97	(0.97, 0.98)	Constant	2.58	(2.60, 2.54)
	0.97	(0.96, 0.99)	R_{t-1}	0.69	(0.59, 0.72)
			\tilde{p}_t	0.29	(0.37, 0.19)
			Frequency	0.29	(0.31, 0.24)
			Inconsistency	0.46	(0.41, 0.53)
			$R_{t-1} \tilde{p}_t$	0.90	(0.70, 1.14)
			R_{t-1} Frequency	0.22	(0.11, 0.29)
		R_{t-1} Inconsistency	0.10	(0.09, 0.30)	

Note: ρ is the daily decay parameter; 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 bootstrapped at the auction level.

D Other results

D.1 Shipper's incentive scheme and instrumental objects

Shippers' incentive scheme. Table 8 presents our estimates of the shipper's incentive scheme. The coefficient of the rejection index R_{t-1} is positive and highly significant, confirming Assumption 1 in our model that shippers punish carriers' rejections with higher probability of demotion. To interpret the magnitude of this coefficient, we simulate two sets of relationships, one with an initial rejection and one with an initial acceptance. We find that a carrier's initial rejection instead of an acceptance reduces the expected number of offers it receives by 3% (from 91 loads to 88 loads). This suggests that the shipper's incentive scheme is soft but generates dynamic incentives that are economically significant.

Consistent with Harris and Nguyen (2021), we find that when the current spot rate (\tilde{p}_t) is high or if the shipper has large volume on a lane, that is, when the relationship is more valuable to the shipper, demotion probability is lower. However, also in such cases the shipper strengthens the incentive scheme, punishing the carrier's rejections more harshly to induce more acceptances.⁴⁶ Additionally, we find that the daily discount rate on the carrier's past rejections is close to one. This suggests that a rejection of the carrier affects the continuation probability of its relationship in many periods.

⁴⁶In Harris and Nguyen (2021), we specified a linear probability model that includes both asset-owners and brokers, and estimate our specification by GMM.

Table 9: The link between spot shares and total market thickness

	Dependent variable $\ln(\text{Volume}^{\text{spot}})$ $\ln(\text{Volume}^{\text{LT}})$							
	5 relationships		10 relationships		15 relationships		20 relationships	
	(OLS2)	(IV)	(OLS2)	(IV)	(OLS2)	(IV)	(OLS2)	(IV)
$\ln(\text{Volume}^{\text{total}})$	0.162 (0.032)	0.257 (0.039)	0.269 (0.037)	0.345 (0.045)	0.272 (0.044)	0.334 (0.052)	0.321 (0.048)	0.408 (0.059)
$\ln(\text{distance})$	0.090 (0.070)	0.001 (0.074)	0.028 (0.075)	0.032 (0.077)	0.002 (0.083)	0.042 (0.085)	0.048 (0.089)	0.097 (0.091)
Frequency	0.261 (0.067)	0.262 (0.067)	0.156 (0.076)	0.159 (0.076)	0.221 (0.109)	0.228 (0.110)	0.366 (0.133)	0.367 (0.134)
Inconsistency	1.597 (0.278)	1.591 (0.279)	1.682 (0.319)	1.653 (0.321)	1.415 (0.429)	1.387 (0.430)	1.325 (0.494)	1.303 (0.497)
origin = MidWest	0.407 (0.082)	0.381 (0.082)	0.378 (0.085)	0.354 (0.086)	0.304 (0.095)	0.281 (0.096)	0.221 (0.098)	0.186 (0.100)
destination = MidWest	0.061 (0.090)	0.030 (0.091)	0.090 (0.096)	0.071 (0.096)	0.141 (0.103)	0.122 (0.104)	0.053 (0.108)	0.029 (0.109)
<u>Instrument</u>								
$\ln(\text{PredictedFlow})$	∅	∅	∅	∅	∅	∅	∅	∅
N	887	887	588	588	427	427	321	321

Notes: Standard errors are in parentheses.

Other instrumental objects. 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 320 days. We estimate that the number of effective bidders is distributed as Binomial(15, 0.21), averaged to 3 effective bidders per auction.

D.2 Robustness of the link between spot market thickness and efficiency

Table 9 presents regression results of Equation (3) as we vary the sets of lane characteristics and the sample restriction to include lanes with at least 5, 10, 15 or 20 relationships (in the microdata). All sample restrictions and specifications show a strong positive link between spot shares and total market thickness.

Table 10 presents our decomposition of the mean transformed costs into the mean operational costs and search costs. Specification (1) is our main specification, where we estimate Equation (17) by two-stage least squares, including as controls a polynomial of degree three of distance and an indicator of market tightness. One potential concern of the main specification is that patterns of trade affect the equilibrium movements of trucks and thus may correlate with unobserved cost shifters. For example, a thick lane may appear desirable not because search

costs are lower on this lane but because it is connected to other thick lanes; this connectivity would help carriers reduce empty miles. To control for such network effects, we construct a measure for the imbalance between forehauls and backhauls, calculated as

$$\text{Imbalance} = \ln(\text{Volume}^{\text{spot}}_{\text{backhaul}}) - \ln(\text{Volume}^{\text{spot}}_{\text{forehaul}}),$$

where $\text{Volume}^{\text{spot}}_{\text{backhaul}}$ denotes the backhaul going from the destination of lane l to the origin of lane l' . This measure captures the likelihood of finding a backhaul, which is the key concern for carriers on long trips (at least 250 miles). Specification (2) includes this measure of volume imbalance on a lane and specification (3) additionally controls for the frequency of interactions and consistency of load timing of the shipper-lane within the contract period. Our estimate of the scale efficiency parameter, the coefficient of $\ln(\text{Volume}^{\text{spot}}) = \text{distance}$, is similar across these three specifications.

Specifications (4), (5) and (6) in Table 10 lend support to our functional form assumption on the relationship between per-mile search costs and spot market thickness. They show that per-mile search costs decrease with the thickness of the spot market, but less so on longer lanes.

D.3 Cluster-specific results

We estimate the distribution of rents and match-specific gains, and perform welfare analysis separately on each of our ten lane-specific clusters and in two market phases.

Figure 18 plots our lane-specific clusters in different shades of gray. The left panel plots the clustered relationships against the average spot rate and log of the average spot volume, which are equilibrium objects used in our K-means clustering. The right panel plots these clusters against their median characteristics, with the sum of search and operational costs representing the supply factor and the log of predicted trade flows representing the demand factor; the size of each cluster represents the number of relationships in that cluster. These scatter plots show that our K-means clustering performs well in separating lanes by their underlying demand and supply factors.


Although the clusters are different in their underlying demand and supply factors, the take-aways from the cluster-specific results are consistent with the aggregated results reported in Section 7. Figure 19 plots the distribution of match quality across all relationships within each lane-specific cluster and each market phase, showing large heterogeneity in match quality within each cluster. Figure 20 plots the median match-specific gains of shippers and carriers in each of our 20 clusters, showing that shippers tend to have larger match-specific gains than carriers from their relationships. Finally, the left panel of Figure 21 shows that under the

Table 10: Cost decomposition

Mean transformed cost	(1)	(2)	(3)	(4)	(5)	(6)
$\ln(\text{Volume}^{\text{spot}})=\text{distance}$	255.850 (77.625)	229.428 (68.241)	229.866 (67.363)			
$\ln(\text{Volume}^{\text{spot}})$				0.883 (0.202)	0.719 (0.151)	0.747 (0.149)
$\ln(\text{Volume}^{\text{spot}}) \ln(\text{distance})$				0.575 (0.441)	0.536 (0.392)	0.518 (0.397)
Tight market	0.536 (0.093)	0.547 (0.086)	0.505 (0.083)	0.548 (0.106)	0.558 (0.092)	0.505 (0.091)
Imbalance		0.359 (0.082)	0.354 (0.080)		0.443 (0.100)	0.446 (0.098)
<u>Instruments</u>						
$\ln(\text{PredictedFlow})=\text{distance}$	∅	∅	∅			
$\ln(\text{PredictedFlow})$				∅	∅	∅
$\ln(\text{PredictedFlow}) \text{ distance}$				∅	∅	∅
<u>Controls</u>						
$h_3(\text{distance})$	∅	∅	∅	∅	∅	∅
Frequency			∅			∅
Inconsistency			∅			∅
N	1162	1162	1162	1162	1162	1162

Notes: Standard errors (in parentheses) are constructed by bootstrapping at the auction level.

Figure 18: Ten lane-specific clusters by K-means method



current fixed-rate contracts, the median relationship in each cluster achieves 40-50% of the relationship-level first-best surplus. The right panel of Figure 21 shows that in the median relationship of each cluster, the shipper gets 75% of the total surplus of the relationship relative to spot transactions.

E Additional figures

Figure 22 confirms that under our estimates of the spot process and shippers' incentive scheme, substantive assumptions on carriers' full compensation and shippers' expected payoff are satisfied. The left panel shows that the carrier's full compensation is increasing in its rent (Assumption 3). The right panel shows that the shipper's expected payoff is increasing in both her rent and the carrier's rent (Assumption 4).

Figure 23 shows a screenshot of DAT load board when a carrier searches for a load. This carrier conducted multiple searches; in each search, it input an origin and a destination, possibly with a radius around these locations, date availability, and some basic information on equipment. The highlighted search is for a load from Houston, Texas to any location on December, 11. The search results show a list of shippers and their contact information. The carrier would contact these shippers and negotiate rates off the platform.

Figure 19: Match-quality (including savings on search costs)

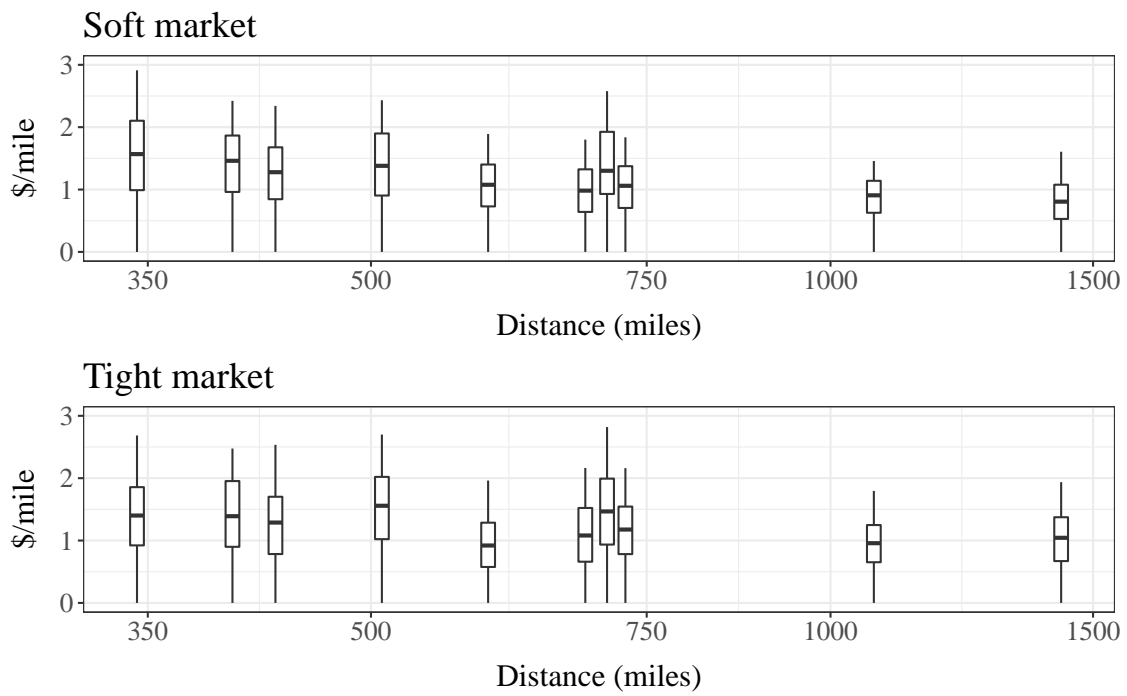


Figure 20: Shippers and carriers' match-specific gains (including savings on search costs)




Figure 21: The median share of total surplus to the first-best surplus and its share by shippers

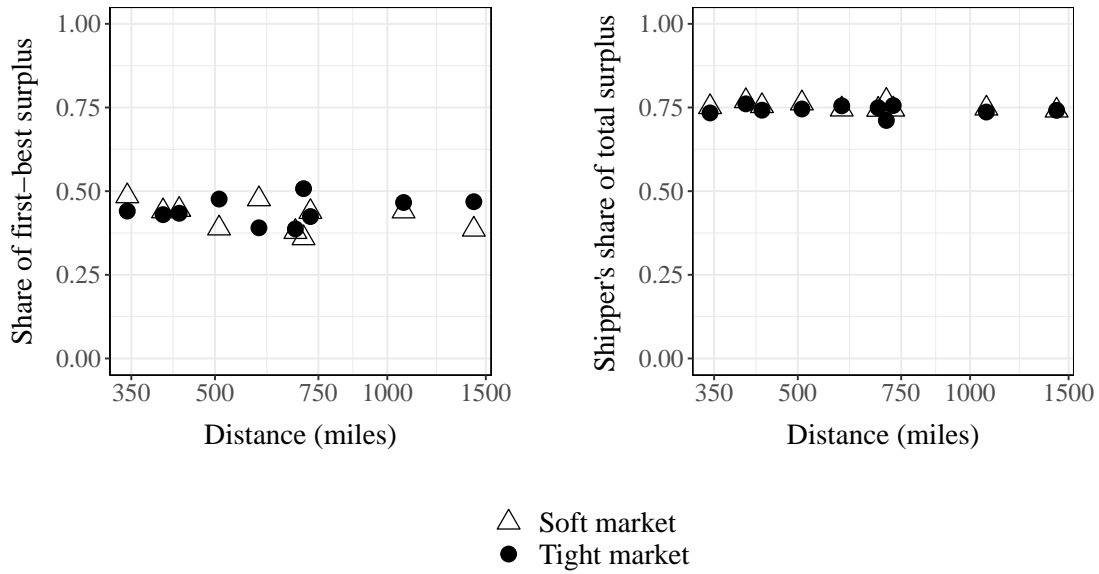


Figure 22: Monotonicity of carrier's full compensation and shipper's expected payoffs in rents

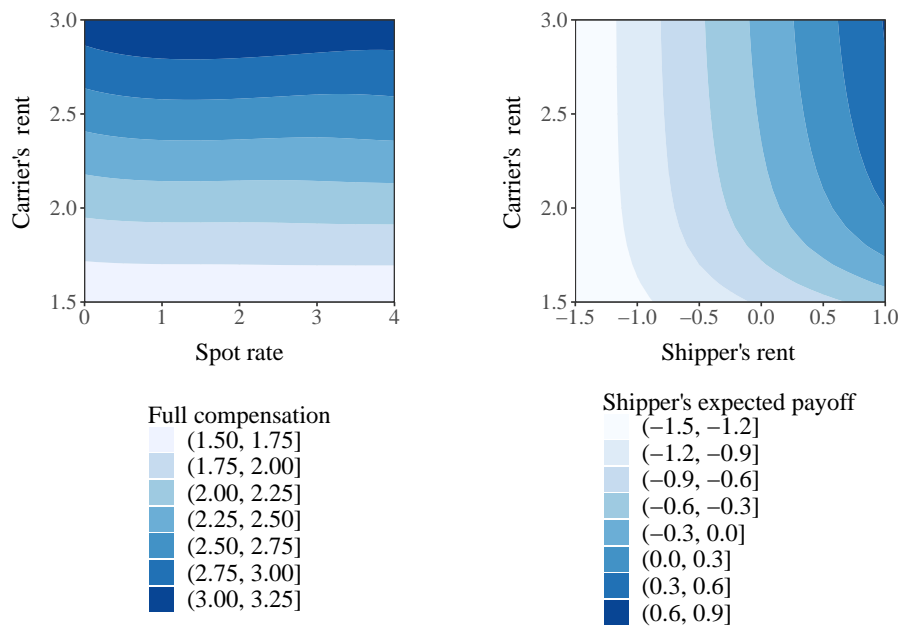


Figure 23: Example of DAT load board

Source: <https://forms.dat.com/resources/product-sheets/dat-load-boards#>.