Pulp Friction: The Value of Quantity Contracts in Decentralized Markets*

Olivier Darmouni
Columbia Business School

Simon Essig Aberg

Juha Tolvanen
University of Vienna

September 4, 2021

Abstract

Firms in decentralized markets often trade using quantity contracts, agreements that specify quantity in advance of trade. We show that firms use quantity contracts to reduce the costs of trading frictions. Specifically, quantity contracts are valuable for two reasons. First, they increase trade between high surplus trading partners because they lock in trade prior to the point of sale. Second, they provide quantity insurance – we show that buyers and sellers are endogenously risk averse with respect to quantity. However, quantity contracts are costly due to their inflexibility to market conditions. Using proprietary invoice data from a large seller, we estimate a model of quantity contracts in the pulp and paper industry. We find that the median value of a quantity contract is 10% of net price. The median value would be 25% lower without quantity insurance and 84% higher without the cost of inflexibility. As trading frictions diminish, the seller uses fewer quantity contracts and profits increase.

Keywords: Decentralized markets, contracts, trading frictions
JEL codes: L14, L22, D23, L73

*Comments welcome. We would like to thank Brian Baisa, Jake Blackwood, Daniel García, Jun Ishii, Cristopher Kingston, Amanda Kowalski, Jessica Reyes, and Charles Smith for helpful feedback.

1 Introduction

Most markets are decentralized and therefore subject to costly trading frictions. For example, search frictions make it costly to find trading partners [Stigler, 1961] and bargaining frictions make it costly to agree on the terms of trade [Rubinstein and Wolinsky, 1985]. How do buyers and sellers structure trade to reduce the costs of trading frictions? The literature has focused on intermediation [Gavazza and Lizzeri, 2021] as a way to increase trade. However, Hsieh and Moretti [2003], Leslie and Sorensen [2014], and Gavazza [2016] show that intermediation is not without costs and may even reduce welfare. In general, alternative structures can arise to mitigate the costs of trading frictions.

In this paper, we show that quantity contracts are valuable in decentralized markets because they allow buyers and sellers to reduce the costs of trading frictions. Quantity contracts are contracts that specify exchanged quantities in advance of trade but sometimes allow prices to be set even after delivery. Many industries use quantity contracts to organize trade, including beef processing [Boyabatlı et al., 2011], coal [Joskow, 1988], and liquefied natural gas [Zahur, 2020]. Our focus will be on quantity contracts in the pulp and paper industry where institutional details allow us to rule out other well-studied reasons firms might use contracts besides reducing the costs of trading frictions. First, contracts do not hedge against price risk [Wolak, 2000] because annual contracts do not lock in a price even though monthly price varies significantly. Second, contracts do not prevent moral hazard on product quality [Lambert, 1983] because pulp is homogenous within well-defined grades. Last, contracts do not reduce incentives to renege opportunistically and 'hold-up' the other party [Williamson, 1979] because contracts often leave price to be negotiated after monthly quantities have been delivered.[1]

Our main contribution is to show that quantity contracts are valuable in decentralized markets for two reasons but at the cost of inflexibility. First, quantity contracts increase trade between high surplus trading partners because they lock in trade prior to the point of sale. We call this the *buyer selection* channel. Trading frictions induce randomness in how buyers and sellers are matched. By signing a contract in advance, a buyer and a seller with high pairwise surplus bypass this randomness and increase expected quantity traded. Concretely, it is logistically difficult for firms to organize trade with many trading partners every month, so firms prefer to sign quantity contracts at the start of the year, locking in trade with high surplus trading partners. More generally, quantity contracts reduce the transaction costs associated with organizing trade. This is also the main benefit of intermediation. In that sense, quantity contracts are a substitute to intermediation, especially in input markets where the identity of buyers and sellers are stable over time.[2]

[1] Sales executives in the industry also reject this view [Brodrechtova, 2015].

[2] In contrast, in markets for assets (such as financial products and business-aircraft), buyers may become sellers and vice

Second, quantity contracts provide *quantity insurance*. Quantity contracts reduce the variance of quantity because they bypass trading frictions that generate variation in quantity. We provide a novel explanation for why reliable, low variance trade is valuable in decentralized markets. We show that when trading frictions generate dispersion in spot prices, buyers and sellers become endogenously risk averse with respect to quantity traded. Consider the perspective of a seller. Due to price dispersion, the first unit sold on the spot receives a higher price than later units. Therefore, the marginal profit of spot trade decreases in total quantity traded on the spot. If the seller allocates its production between a contract buyer and spot buyers, then the decreasing marginal profits from the spot market imply that the seller's total profits are concave with respect to the quantity it trades with the contract buyer. Thus, the seller prefers for the trade with the contract buyer to have low variance. Concretely, contract buyer reliability diminishes the risk that a seller must trade with relatively undesirable spot buyers.

The main cost of quantity contracts is *inflexibility* to market conditions. Quantity contracts lock in the distribution of trade before market conditions realize and prevent buyers and sellers from reallocating quantity ex-post. Concretely, if a seller unexpectedly comes into contact with desirable spot buyers some month, then quantity contracts may prevent the seller from trading with these buyers if production is inelastic. One reason that annual contracts may not have a longer term structure is to increase flexibility to market conditions. Industries where flexibility is more important may not rely on quantity contracts.

Our second contribution is to develop an empirical model of quantity contracts and spot trade and use it to quantify the value and cost of quantity contracts in the pulp industry. We use proprietary invoice-level data from a large seller in the pulp industry from 2014 to 2019. In addition to the price and quantity of each transaction, we observe contract fulfillment rebates, logistics costs, and production costs. We observe invoices for both contract and spot trade. We combine these data with publicly available information on market conditions such as regional gross price indices and average rebates. The detailed breakdown of costs, prices and rebates gives us a clear picture of profit margins, expected pairwise quantities, and probabilities of trading, while the global variables allow us to control for market-level shocks.

These microdata are critical for quantifying the value and cost of quantity contracts. Specifically, we need to observe the full distribution of profit margins in contract and spot trade to measure the extent to which quantity contracts match our seller with the highest surplus buyers. Similarly, we use the full distribution of spot prices within a month and the full distribution of contracted quantities across time to quantify the value of quantity insurance. The high frequency data on spot trades allows us to quantify the costs of inflexibility. A limitation of our data is that we only observe one seller, so our focus is on bilateral contracting instead of industry-wide equilibrium.

versa.

We first use this microdata to establish stylized facts that point to the value and cost of quantity contracts. Conceptually, the value of contracts relies on two key assumptions, and we find support for these assumptions in the data. First, there is spot price dispersion. The standard deviation of spot prices is 11% of the mean within a market and month, consistent with trading frictions. Spot buyers trade with lower probability than contract buyers regardless of price, also consistent with trading frictions. Second, we find reduced-form evidence that the seller's total sales are inelastic to market price in the short run. Institutional details support this finding; production is highly capital intensive and scaling it down is costly due to technological constraints. This inelasticity implies that the seller allocates its production between contract buyers and spot buyers, generating concavity in its total profit function.

We also find reduced-form evidence in the data consistent with quantity contracts being valuable due to buyer selection and quantity insurance but costly due to inflexibility. First, we find that contract buyers have lower logistics costs than spot buyers and that average contract buyer mill gate prices are higher than average spot buyer prices (mill gate price equals gross price minus rebates minus logistics costs). Second, we find that the average contract buyer coefficient of variation in quantity is 61% lower than that of spot buyers. Third, in 40% of months, the minimum contract price is lower than the minimum spot price, so the seller would have ex-post preferred to allocate more quantity to spot buyers.

Informed by the stylized facts, we develop an empirical model of quantity contracts and spot trade. In the model, the seller has inelastic production that it allocates between spot buyers and contract buyers. On the spot, the seller only comes into contact with a fraction of the available buyers each month due to trading frictions. This fraction indexes the extent of trading frictions. When trading frictions diminish, the seller comes into contact with more available buyers. The seller faces a constraint on how much it can trade to spot buyers, so it must be selective about the buyers with whom it trades. The spot buyers have heterogeneous outside options, generating price dispersion. The seller makes take-it-or-leave-it offers to the most desirable buyers among the set of buyers that it contacts.

We model quantity contracts based on the structure of contracts in the industry. Quantity contracts commit the buyer and seller to trade a stable monthly quantity. They also commit the buyer and seller to a rebate off of market price. Market price realizes after the contract is signed. Contract buyers bargain over the rebate through a Nash-in-Nash bargaining protocol [as in Crawford and Yurukoglu, 2012, Grennan, 2013, Collard-Wexler et al., 2019]. If bargaining fails, then the contract buyer joins the pool of spot buyers.

We estimate the model parameters in order to quantify the value and cost of quantity contracts. The rich data facilitate transparent identification of the model parameters. To identify the distribution of outside options, we use the observed prices and quantities of spot trades. To identify the fraction of available spot buyers that the seller comes into contact with each month, we use the correlation between spot buyer trade probabilities and spot buyer price. If the most desirable spot buyers trade with the same probability as less desirable spot buyers, then the fraction must be low because the seller is not selective among spot buyers. In contrast, if high price spot buyers trade with greater probability, then the seller is somewhat selective. Finally, to identify contract buyer bargaining powers, we compare contract buyer rebates to comparable spot buyer rebates. The larger the contract buyer rebate, the greater the buyer's relative bargaining power.

We estimate the model in a two step procedure. The first step uses generalized method of moments on the sample of spot buyers to estimate the distribution of buyer outside options and the fraction of available spot buyers that the seller comes into contact with each month. The second step uses regressions on the sample of contract buyers to estimate other parameters of the buyer outside option distribution and the contract buyer relative bargaining powers. We find that the seller comes into contact with 78% of available spot buyers, allowing the seller to be somewhat selective, but not perfectly selective. For context, the average spot buyer probability of trade is 65%. We estimate a median contract buyer relative bargaining power of 38%. In support of this low buyer bargaining power, buyers tend to be smaller than sellers in the industry.

Using the estimated parameters, we simulate three counterfactuals. The first counterfactual quantifies the value of quantity contracts and establishes comparative statics. We find that the median quantity contract is worth €57/ton, or 10% of average net price. Consistent with the buyer selection channel, the value of contracts is larger when contract buyers have lower logistics costs. Consistent with the quantity insurance channel, the value decreases when contract buyers have larger variance of quantity. Consistent with the inflexibility channel, the value increases with the left truncation on the distribution of spot buyer rebates. We then find that the value of quantity contracts decreases when trading frictions diminish. If trading frictions are minimal, in that the seller always comes into contact with all the spot buyers we observe in our data, only 20% of quantity contracts are valuable. This result is consistent with the primary purpose of quantity contracts being to reduce the costs of trading frictions.

The second counterfactual isolates the value due to quantity insurance and the cost due to inflexibility. To eliminate the value due to quantity insurance, we turn to the theoretical comparative static that the value of contracts decreases in the variance of contracted quantity. We replace each contract buyer's variance of quantity with the variance of their quantity if that buyer were to trade on the spot. After this replacement, the median value of a quantity contract falls to €43/ton, a 25% decrease from the baseline. To eliminate the cost of inflexibility, we allow the seller to renege on contracted quantities ex-post. The median value of a quantity contract increases to €105/ton, an 84% increase from the baseline.

The third counterfactual evaluates how contracting behavior changes with the magnitude of trading frictions. We assume that the seller signs quantity contracts to maximize total profits. We approximate the profit-maximizing outcome by specifying an algorithm that allocates buyers between contract and spot.

The algorithm predicts that 84.9% of the seller's total sales are contracted under the baseline parameter estimates. For comparison, 84.6% of the seller's total sales are contracted in the data. We find that when trading frictions are minimal, only 30% of total quantity is contracted.

Trading frictions and quantity contracts have large implications for welfare. Trading frictions decrease the seller's profits by 9% relative to a baseline where trading frictions are minimal. In contrast, if the seller is forbidden from signing quantity contracts, then its profits are 14% below the baseline. Therefore, quantity contracts meaningfully reduce the cost of trading frictions.

1.1 Related Literature

First, we contribute to the literature on trading frictions in decentralized markets. Our main contribution is to the literature that estimates how and to what extent certain structures such as intermediaries avoid the cost of trading frictions [Gavazza, 2016, Egan, 2019, Farboodi et al., 2018, Donna et al., 2019].[3] To our knowledge, we are the first to study how quantity contracts serve this purpose. Gavazza [2011] finds that leasing reduces transaction costs in the market for commercial aircraft, and thus serves a similar role as intermediation. Instead of studying an asset market where the identity of buyers and sellers may change, we study an input market where these roles are stable. Methodologically, we are most similar to Salz [2020], who quantifies search costs and finds that intermediaries affect welfare both directly and indirectly in the market for trade-waste. Instead of attempting to quantify the distribution of search costs, we use our rich data to incorporate heterogeneity into buyer outside options, and search costs may be one component of these outside options. We do not attempt to separately estimate search costs because we only observe one seller and thus do not observe buyer search behavior, but we are still able to quantify the welfare effects of quantity contracts. Furthermore, several empirical papers document trading frictions as the cause of price dispersion in decentralized markets [Chandra and Tappata, 2011, Kaplan and Menzio, 2015, Stango and Zinman, 2016], consistent with the price dispersion we document in our data.

Second, we contribute to the literature that studies contracting relationships between firms by providing a novel rationale for the value of contracts and developing an empirical model of contracts in decentralized markets. A large literature studies how the structure of contracts responds to frictions, both in theory [Aghion and Bolton, 1992, Dewatripont and Tirole, 1994, Bajari and Tadelis, 2001, Strulovici, 2017] and empirically [Joskow, 1987, Corts and Singh, 2004, Vanneste and Puranam, 2010].[4] Most related to our work, Zahur [2020] finds that long-term contracts prevent hold-up on capital investment in the market for liquefied natural gas, and estimates a structural model of contracts and spot trade to quantify these effects. We focus

[3] See Gavazza and Lizzeri [2021] for a review.
[4] See Lafontaine and Slade [2010] for a review.

on the value of contracting in a setting where hold-up is less a concern, and thus provide a complementary explanation for the value of contracts. The main cost of contracts, inflexibility, is similar in the two contexts. Extending the model, we incorporate trading frictions into our model of spot trade, which are crucial to understanding the value of contracts in our setting. Another strand of this literature studies relational contracts that rely on dynamic enforcement [Baker et al., 2002, Halac, 2012, Malcomson, 2012, Gil and Marion, 2013, Li and Matouschek, 2013, Helper and Henderson, 2014, Fong and Li, 2017, Gil and Zanarone, 2018]. Most related, Macchiavello and Morjaria [2015] find that buyer beliefs about seller reliability are important in the Kenyan rose export sector. In that paper, the value of reliability is given exogenously by a kink in the payoff function and their focus is on the dynamic evolution of the value of relationships due to learning about the seller's reliability. We complement this approach by showing why reliability may be endogenously valuable in decentralized markets, even for risk-neutral firms. The relationships we observe are stable and mature, so learning about reliability is likely to play a smaller role. Furthermore, contracts in the pulp industry include so-called 'contract fulfillment rebates' that lend some enforceability to contracts and induce reliability.

2 Background and Data

In this section, we provide background on the pulp and paper industry, describe the organization of trade, and introduce the data.

2.1 The Pulp and Paper Industry

Fiber is the most important input in the production of paper products including tissue, printing and writing paper, specialty papers, and packaging materials. 430 million tons of fiber are produced annually, of which 60% are recycled materials and 40% are pulp.[5] Pulp and recycled materials are not substitutable because most paper mills have machinery specific to one of the two. Within the pulp industry, there are two methods for producing pulp from wood (each accounting for about half of pulp production): mechanical and chemical. Paper mills usually use only one of these two types due to machine specificity and differences in end-use. We restrict analysis to the chemical pulp industry, with annual global production of 105 million tons, and trade valued at approximately €60 billion per year.[6]

Seller production is relatively inelastic to market conditions in the short run. This is true because pulp mills face high fixed operational costs, so mills tend to produce as close to capacity as possible. In the mid to

[5]The statistics in this section are from internal presentations by one large seller in the industry.
[6]Specifically, we consider bleached chemical pulp, which is by far the largest segment of chemical pulp.

long run, production responds to market conditions through capacity expansions and the construction of new mills. Nevertheless, there is still monthly variation in production due to planned maintenance, unplanned breakages, and worker strikes. Variation in seller production is one force that causes variation in a buyer's willingness to pay and residual demand over time. In response to the inelasticity of production in the short run, pulp producers use inventory. However, inventory is constrained by storage limitations and the opportunity cost of delaying sales.

Buyer consumption is also inelastic to market conditions in the short run. Like pulp mills, paper mills face high fixed operational costs, so mills tend to consume as close to capacity as possible. Buyers tend to be smaller than sellers, and may not be able to inventory significant quantities of pulp. As such, they face a large opportunity cost of running below capacity.

Up to a few well-defined characteristics, pulp is homogenous. There are two types of pulp: pulp from hardwood (70%) and pulp from softwood (30%). To further subdivide these categories, there are a few different types of hardwood and softwood depending on the species of the source tree. Conditional on tree species and bleaching method, pulp is homogenous. Invoices often include a standardized document with technical specifications. Given the homogeneity of pulp, one might expect to see a large and deep spot market because moral hazard on product quality is not a major concern. In reality, most trade occurs via quantity contracts and the spot market is thin. The goal of this paper is to explain and quantify the value of these contracts in decentralized markets for homogeneous goods.

2.2 Organization of Trade

Structure 1: Quantity Contracts. Most trade in the pulp industry occurs through annual quantity contracts. These contracts specify an annual quantity target. The contracts also specify that quantity should be stable from month to month. Sometimes, buyers purchase multiple fibers from a seller and have multiple quantity targets within a single contract.

Gross prices are either indexed or negotiated each month. In most contracts, the gross price is indexed. There is no single market price, but a few consulting firms survey large buyers and sellers and release price indices each week or month for each fiber and region. A typical indexed contract specifies a past price index or average of past price indices to determine the gross price. In other contracts, and especially contracts with large quantity targets, the gross price is negotiated each month. A representative of one large seller in the industry suggested that gross prices are negotiated in order to aggregate information on market conditions, and monthly price negotiation does not reflect changes in relative bargaining positions.

Regardless of whether gross price is indexed or negotiated, quantity contracts include a base rebate off this gross price. Differences in the rebate across contracts reflect differences in relative bargaining positions. Since these rebates are negotiated once annually, the split of the surplus is relatively stable within a year.

Quantity contracts have strict enforcement mechanisms. The main enforcement concern is that the market will change unexpectedly and the buyer will refuse to trade with the seller or vice versa. To alleviate this concern, contracts include a 'contract fulfillment rebate.' This clause specifies an additional rebate that is paid if and only if the terms of the contract are satisfied at the end of the year. These rebates range from 0.5% to 6% of gross price, but the average is less than 1.5%. Larger contracts include more complicated contract fulfillment rebate structures, such as multiple rebates for achieving various quantity targets. One seller noted that buyers with performance rebates written into an annual contract almost always receive them because the bonus is sufficiently large.[7] In other words, it seems that these contract features are sufficient to prevent opportunistic behavior for contract buyers. Dynamic considerations serve as a secondary enforcement mechanism. Many buyers and sellers have relationships dating back to the 1980s or earlier, so preserving reputation may be important. Given these enforcement mechanisms, we abstract from enforcement concerns in our model and assume that contracts are enforceable for contract buyers.

Structure 2: Spot Trade. Trade that does not occur via quantity contracts occurs on the spot. Price and quantity are negotiated shortly before the transaction occurs. These transactions tend to be smaller and usually occur only a few times a year between a given buyer and seller. Many buyers and sellers trade via quantity contracts with some trading partners and trade on the spot with others.

Spot trade is subject to trading frictions. During our analysis period, there was no centralized exchange where spot trade occurred. Two trading frictions may inhibit efficient spot trade. First, it is costly to search for spot buyers because doing so requires engaging with sales managers. Second, it is costly to bargain over price and quantity each month. Specifically, variation in production, demand, and market conditions make it costly to determine the price and quantity a spot trading partner is willing to accept. Since 2020 (after our analysis period), a large physical futures exchange has developed in Shanghai, though a number of exchanges failed in the preceding decade due to insufficient volume. One function of this futures exchange may be to facilitate the matching of buyers and sellers, thus alleviating trading frictions.

Structure 3: Vertical Integration. Of the ten largest pulp buyers, seven are partially vertically integrated.[8] This wave of vertical integration, especially among European firms, largely occurred in the 1980s through the acquisition of pulp mills by paper mills. All of the largest vertically integrated firms trade a significant portion of pulp externally instead of relying on internal pulp transfers. Wang [2005] provides evidence that vertical integration occurred most prominently in market segments with the highest

[7]In our data, we are unable to observe cases where the contract fulfillment rebate was not paid.

[8]Several papers analyze the causes [Ohanian, 1994, Niquidet and O'Kelly, 2010, Kimmich and Fischbacher, 2016] and effects [Pesendorfer, 2003] of integration in the industry.

concentration, suggesting that the threat of opportunistic behavior motivated integration. Even if integrated buyers and sellers choose to trade most pulp externally, the option value of trading internally is sufficient to dissuade external trading partners from opportunistically renegotiating contracts.

2.3 Data

We analyze proprietary invoice data from a large seller in the pulp industry from 2014 to 2019. Each monthly invoice documents a pulp transfer from one of the seller's pulp mills to a paper mill for a particular pulp product. Each invoice contains information about the location of delivery, product, month, terms of payment, price, rebates, logistics costs, storage costs, and variable production costs. We merge these invoices with the production, inventory, and delivery data from each of the seller's pulp mills, as well as with an internal estimate of market price and average rebate. We focus on transactions between non-integrated mills.

We merge these invoices with the seller's internal buyer classification system. The seller classifies buyers into four tiers. Tiers One and Two are primarily comprised of contract buyers, and Tiers Three and Four are primarily comprised of spot buyers. Among the first two tiers, the seller says it prioritizes Tier One buyers because those buyers are most important to its business. Among the second two tiers, the seller says it prioritizes Tier Three buyers because those buyers have high willingness to pay or low logistics costs. We therefore say that a buyer is 'prioritized' if it is classified in Tier One or Tier Three.

We categorize buyers into contract and spot using the contract fulfillment rebate and the internal buyer classification system. We classify the buyer as a contract buyer if it has a positive contract fulfillment rebate in the invoice data or it is in Tier One. We allow the classification to change from year to year. Furthermore, because some buyers operate in multiple regions and purchase multiple fibers, we allow the classification to vary by region and fiber. In the data, the seller services 268 buyers, of which 75 are contract buyers and 193 are spot buyers. Despite this imbalance, the contract buyers account for more than 83% of all quantity traded.

Table 1 provides unweighted descriptive statistics of the buyer characteristics, invoices, seller variables, and market variables. For some variables, the mean and median are removed to protect the seller's anonymity. Note that the seller services a global portfolio of buyers (the region identities are obscured for the sake of anonymity). Buyers are heterogeneous in size. The mean buyer capacity is far larger than the median, though we only observe capacity for a subset of buyers. The median buyer only traded in 15 of the 72 months in the sample.

Because our categorization of buyers varies over fibers, regions, and years, our main unit of analysis throughout is a buyer that operates within a particular fiber, region, and year. Therefore, we refer to a

Table 1: Descriptive Statistics

	N	Mean	SD/Mean	Median	IQR
Panel A: Buyer characteristics.					
Contract buyer	268	0.28			
Fiber: Hardwood	268	0.63			
Fiber: Softwood	268	0.37			
Region: A	268	0.66			
Region: B	268	0.34			
Internally prioritized	268	0.29			
Capacity (tons/month)	136	24,101	1.89	9,186	16,209
Months with positive trade	268	24.22	0.97	15	32
Panel B: Invoices.					
Quantity (tons)	6,492		1.57		2,236
Total rebate (% gross price)	6,492		0.38		0.15
Contract fulfillment rebate (% gross price)	6,492		1.66		0.01
Logistics costs (€/ton)	6,492		0.36		26.65
Mill gate price (€)	6,492		0.19		124.37
Production costs (€/ton)	6,492		0.13		44.04
Panel C: Seller variables.					
Production (tons/month)	72		0.09		32,133
Inventory (tons/month)	72		0.12		50,579
Total Sales (tons/month)	72		0.07		24,892
Panel D: Market variables.					
Market price (China, hardwood, €)	72	645.14	0.16	630	140
Market price (Europe, hardwood, €)	72	817.27	0.16	775	217.5
Market price (China, softwood, €)	72	700.49	0.15	672.5	125
Market price (Europe, softwood, €)	72	932.67	0.15	897.50	180
Market rebate (China, hardwood, % gross price)	6		0.03		0.01
Market rebate (Europe, hardwood, % gross price)	6		0.1		0.04
Market rebate (China, softwood, % gross price)	6		0.04		0
Market rebate (Europe, softwood, % gross price)	6		0.12		0.05

Notes. Certain statistics are excluded to preserve the anonymity of the data provider, a large pulp seller. The data span from 2014 through 2019. Some buyers purchase multiple fibers and operate in multiple regions, so those statistics in Panel A are quantity-weighted averages among the buyers. Capacity is unavailable for some buyers. Invoices are at the buyer-fiber-region-month level. All numeric variables are winsorized at the 0.1% level. All price and cost variables are in January 2015 Euros, and all quantity variables are in tons. Logistics costs are the difference between price after rebates and mill gate price. Seller variables are at the month level. Market price is at the month level and average rebate is at the annual level.

buyer within a fiber, region, and year as a 'buyer' to facilitate exposition. We refer to a given fiber and region as a 'market.'

Panel B shows the elements of an invoice. Starting with a gross price, the seller deducts a base rebate and a contract fulfillment rebate to arrive at price after rebates. Then, subtracting off logistics costs, we arrive at the mill gate price. After subtracting production costs, we have the sales margin. One representative of the seller noted that the seller aims to maximize trade at a high mill gate price, instead of primarily considering the sales margin. This is true because the seller considers production costs as largely fixed due to the short-run inelasticity of production. Therefore, we use mill gate price as the primary price measure.

Panel C shows the available seller variables and Panel D the available market variables. Market price is a measure of gross price, not net price. There is large variation in market price across regions and fibers, but much of this difference can be explained by corresponding differences in market rebate. Rebates range from under 5% to over 25% across regions and fibers. Differences in price after rebate across regions and fibers are much smaller.

3 Stylized Facts

In this section, we present five stylized facts. The first and second facts support two key assumptions of the model. The third, fourth, and fifth facts point to the value and costs of quantity contracts in terms of buyer selection, quantity insurance, and inflexibility.

Key Assumption 1: Spot Price Dispersion. The first key assumption of the model is that there is spot price dispersion, consistent with trading frictions. Figure 1 presents a histogram of mill gate price within a region, fiber, and month relative to the average price. Within a typical month, prices range from 20% below the mean to 20% above the mean and the standard deviation in price is 11% of the mean.

Figure 1: There Is Substantial Spot Price Dispersion

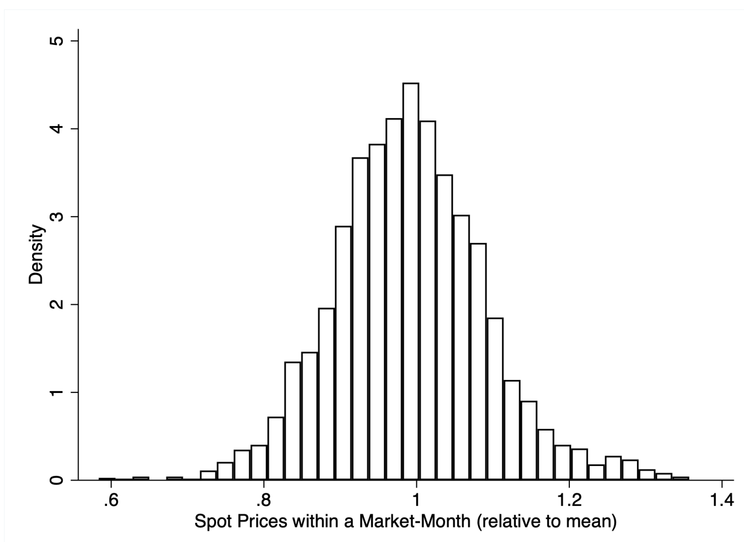

Notes. Mill gate price among spot buyers after removal of a market-month fixed effect, relative to mean of one. Market is defined as product-region.

In some models of trading frictions, price dispersion arises as a result of mixed strategy pricing. In others, price dispersion arises because trading partners have heterogeneous outside options. We find evidence consistent with the latter mechanism. Figure 2 presents a binscatter of buyer trade probability on the average mill gate price that the buyer pays. Among spot buyers, the probability of trade increases in average mill gate price. Table OA.1 shows that this trend is significant. This trend is consistent with a model where buyers have heterogeneous outside options and the seller is somewhat selective when choosing with whom it trades. An implication is that the seller faces a downward-sloping demand curve on the spot, in that the seller receives a higher price for the first units it sells on the spot. Figure OA.1 provides some examples of these demand curves. A secondary takeaway of Figure 2 is that spot buyers trade at lower probability than contract buyers, even when controlling for the average mill gate price. Table OA.1 shows that this difference is significant. The lower trade probability of spot buyers is consistent with a model where trading frictions prevent a buyer and seller from meeting with some non-zero probability.

Figure 2: More Desirable Spot Buyers Are More Likely to Trade

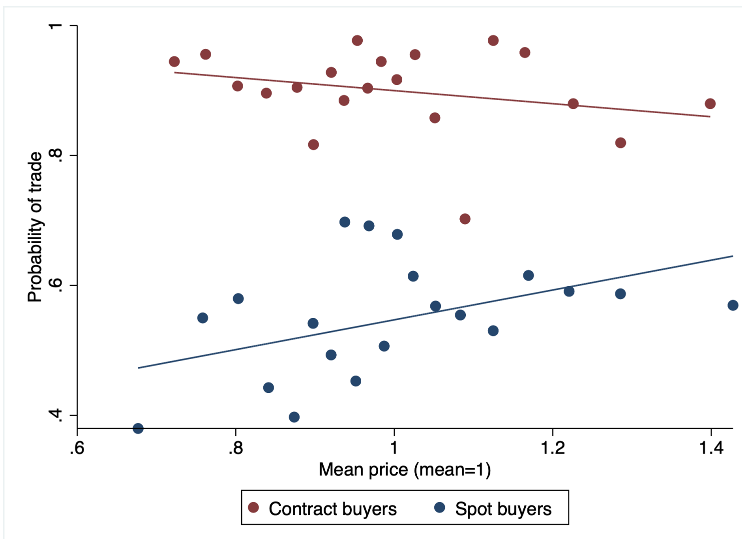

Notes. Each observation is a buyer within a particular region, fiber, and year. The price measure is average mill-gate price paid within a year, relative to the full-sample mean. Observations are binned.

Key Assumption 2: Inelastic Total Sales in the Short Run. The second key assumption of the model is that the seller's total sales are inelastic in the short run. The first column of Table 2 shows that the reduced-form elasticity of total sales within a year with respect to market price is statistically insignificant and quantitatively small. To decompose where the reduced-form elasticity into its components, Table 2 shows that the reduced-form elasticity of production within a year is nearly zero, but that the reduced-form elasticity of inventory is somewhat larger and significant. These results correspond to the institutional details that pulp mills always aim to run at or near capacity, but there is some ability to use inventory. Because inventory is relatively inelastic to market conditions, for simplicity the baseline model abstracts from it.

Value of Quantity Contracts 1: Buyer Selection. Quantity contracts are valuable if they increase trade with the highest surplus trading partners. A major component of the seller's costs are the logistics

Table 2: Total Sales Are Inelastic to Market Price in the Short Run

	(1) Log Total Sales	(2) Log Production	(3) Log Inventory
Log Market Price	-0.17	-0.01	-0.38
	(0.11)	(0.16)	(0.15)
Observations	72	72	72
R-squared	0.19	0.07	0.28
Year FE	YES	YES	YES

Notes. Robust standard errors in parentheses. Observations are at the monthly level. Market price is the average market price among the seller's trading partners.

Table 3: Contract Buyers Have Lower Logistics Costs

	(1)	(2)	(3)
	Logistics Costs		
Contract Buyer	-6.94	-8.31	-8.51
	(1.58)	(1.45)	(1.89)
Capacity (million tons)			0.95
			(9.88)
Prioritized			-6.25
			(1.95)
Observations	797	797	590
R-squared	0.03	0.19	0.24
Market-year FE	NO	YES	YES

Notes. Robust standard errors in parentheses. Each observation is a buyer within a region, fiber, and year. Logistics costs are measured in €/ton.

costs of transportation from the pulp mill to the paper mill. Therefore, paper mills that are located near the seller are more desirable than other mills. Table 3 provides evidence that contract buyers have lower logistics costs than spot buyers. As further evidence that quantity contracts increase trade with desirable trading partners, mill gate price is higher on average for contract buyers. Figure 3 shows the ratio of average contract mill gate price and average spot mill gate price within a region, fiber, and month. Ex-ante, the average contract price is 7% higher than the average spot price, and this is statistically significant.

Value of Quantity Contracts 2: Quantity Insurance. Quantity contracts serve as a form of quantity insurance if they reduce the volatility of trade. Figure 4 plots the coefficient of variation in quantity (including zeros when trade does not occur) for contract buyers and spot buyers. The average contract buyer coefficient of variation is 0.53 and the average spot buyer coefficient of variation is 1.38. In Table OA.2, we show that this difference is statistically significant, even conditional on other observable buyer characteristics. We use the coefficient of variation instead of variance to account for differences in scale between contract and spot buyers. Much of the difference in volatility comes from the fact that spot buyers tend to trade less

Figure 3: Average Contract Prices Are Higher Ex-Ante

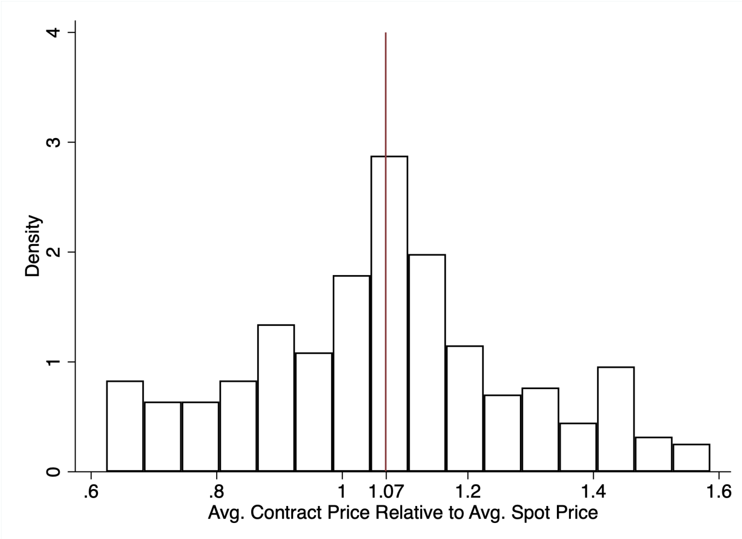

Notes. The median and mean values equal 1.07. Each observation is trade within a region, fiber, and month. The price measure is mill gate price.

frequently than contract buyers. In Figure OA.2 and Table OA.2, we show that the difference between the coefficients of variation in quantity conditional on trade (that is, excluding the zeros) is about one-tenth as large. Therefore, in our model, we assume that the variance of quantity conditional on trade is the same regardless of whether a buyer trades through a contract or on the spot.

Cost of Quantity Contracts: Inflexibility. The main cost of quantity contracts is that they are inflexible to changing market conditions. Figure 5 shows the ratio of minimum contract mill gate price and minimum spot mill gate price within a region, fiber, and month. In 40% of cases, the minimum contract price is below the minimum spot price. In these cases, the seller would have preferred ex-post to allocate more quantity to spot buyers and away from contract buyers.

4 The Value of Quantity Contracts: The Basic Idea

In this section, we describe a simplified model of quantity contracts and spot trade. The purpose of the simplified model is to clarify how quantity contracts generate surplus through buyer selection and quantity insurance but come at the cost of inflexibility. The most important simplifications of the model are that it assumes spot buyers are infinitesimal, and only considers one potential contract buyer instead of many. In Section 5, we present the full empirical model that relaxes these and other assumptions.

We model trade as a game between a single large seller and many smaller buyers. We do not model aggregate market conditions because we only observe data from a single seller. We suppose that the seller

Figure 4: Contract Buyers Are More Reliable Than Spot Buyers

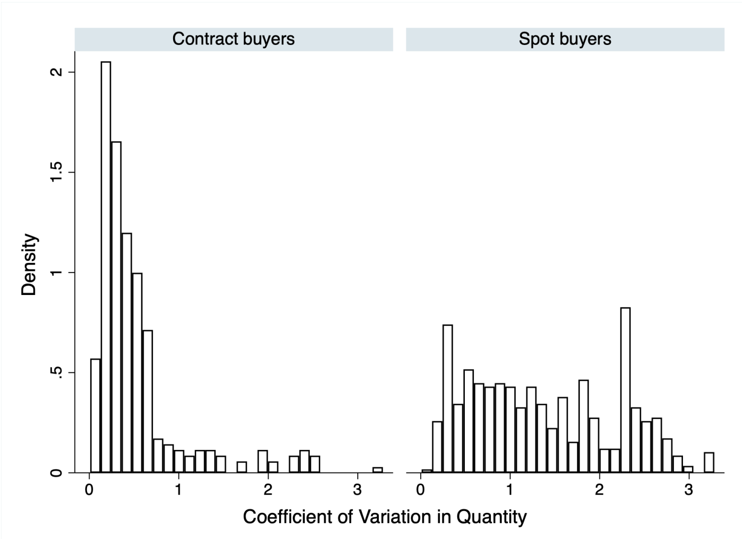

Notes. The coefficient of variation equals the standard deviation of quantity over the mean quantity. Each observation is a buyer within a region, fiber, and year.

Figure 5: Minimum Contract Prices Are Sometimes Lower Ex-Post

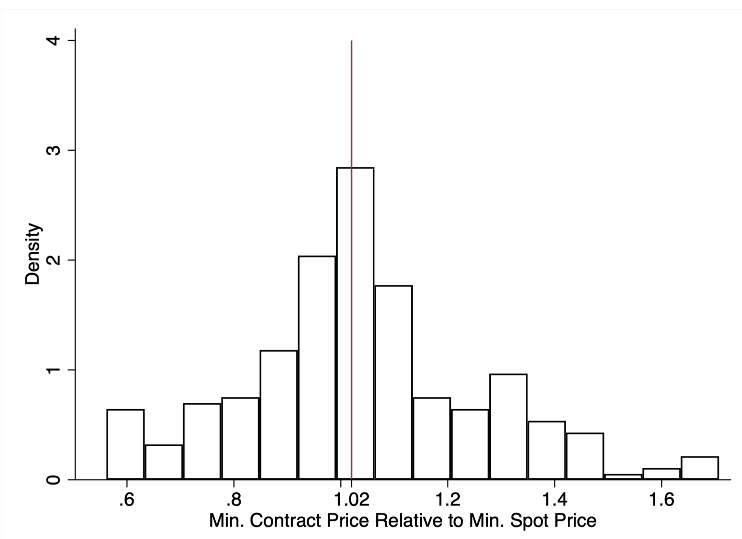

Notes. The median value equals 1.02 and the mean value equals 1.04. For the purposes of this figure, prices are winsorized at the 5% level. Each observation is trade within a region, fiber, and month. The price measure is mill gate price.

decides whether to sign a contract with buyer i. This buyer has willingness to pay p_i for all units traded. In the simplified model, we assume the seller can perfectly price discriminate, so it charges a price p_i whenever it trades with buyer i. We assume that buyer i can only trade zero or Y_i units with the seller.

We make two key assumptions, as we motivate in Section 3. First, we assume there is spot price dispersion. We model the spot market as a set of infinitesimally small buyers with variable willingness to pay. The seller comes into contact with a subset of these small spot buyers and trades with those buyers that have the highest willingness to pay. The $\tilde{q}$th unit the seller trades on the spot has price $p_S(\tilde{q})$. The marginal spot price curve $p_S(\tilde{q})$ decreases in $\tilde{q}$. The top left panel of Figure 6 shows an example marginal spot price curve $p_S(\tilde{q})$. Second, we assume inelastic total sales. That is, the seller has a fixed production equal to Q, and it must allocate this production between buyer i and the spot buyers. If the seller trades q_i with buyer i, then it trades the remaining $Q - q_i$ units with the spot buyers and earns spot revenue equal to the area under $p_S(\tilde{q})$ for $\tilde{q} \in [0, Q - q_i]$.

The top right panel of Figure 6 plots the seller's profits as a function of the quantity that it trades with buyer i q_i, under the assumption of constant marginal cost κ. To start, we suppose that the spot price curve $p_S(\tilde{q})$ is deterministic, abstracting away from volatility in market conditions that generates the costs of inflexibility. When the seller does not trade with buyer i and q_i equals zero, then the seller's revenue is simply the area under the marginal spot price curve until it exhausts its production Q. When the seller trades with buyer i and q_i equals Y_i, then the seller receives revenue $p_i Y_i$ from buyer i, and trades the remaining $Q - Y_i$ units to the small spot buyers according to the marginal spot price curve.

We consider two ways that the seller could structure trade with buyer i. If buyer i trades on the spot, trading frictions restrict contact between buyer i and the seller such that buyer i comes into contact with the seller with probability $\gamma \in (0, 1)$. If the seller signs a quantity contract with buyer i, then trade with buyer i always equals Y_i. That is, the quantity contract allows the buyer and seller to bypass the trading friction.

The seller's expected profits when it signs a quantity contract $\mathbb{E}[\Pi^C]$ is higher than the seller's expected profits when buyer i trades on the spot $\mathbb{E}[\Pi^S]$. Therefore, quantity contracts are valuable. This is true for two reasons:

- **Buyer Selection.** A quantity contract is valuable because it allows the seller to bypass the trading friction and trade more with high willingness-to-pay buyer i. When the seller writes a quantity contract, the expected trade with buyer i increases from γY_i to Y_i. Because buyer i has a high willingness to pay p_i, an increase in the expected trade with buyer i increases the seller's expected profits. Due to the second key assumption of inelastic total sales, expected trade with low willingness-to-pay spot buyers

Figure 6: The Value of Quantity Contracts

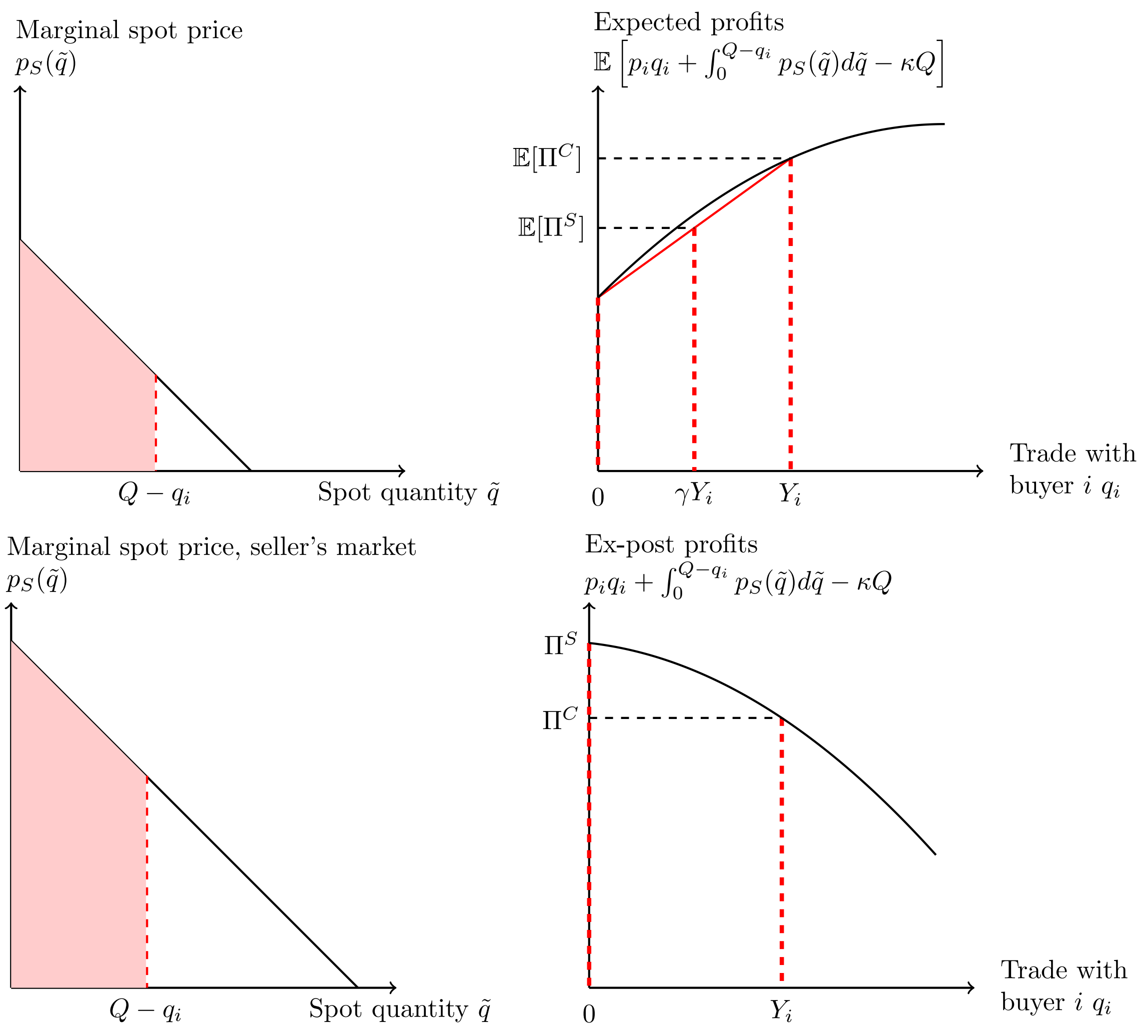

decreases when the seller signs a quantity contract with buyer i. If buyer i's willingness to pay p_i is greater than the average of the marginal spot price curve $p_S(\tilde{q})$ over $[Q - Y_i, Q]$, the quantity contract is valuable.

- **Quantity Insurance.** A quantity contract is valuable because it reduces the volatility of quantity. Variable quantities expose the seller to the risk of trading with low willingness-to-pay spot buyers. When the seller writes a quantity contract, the variance in quantity traded with buyer i decreases from $\gamma(1-\gamma)Y_i^2$ to zero. The second panel of Figure 6 shows that the seller's profit is concave in quantity traded with buyer i. This is true regardless of how we plot the marginal spot price curve $p_S(\tilde{q})$ as long as the curve is decreasing. Formally, because the seller's profits from trading q_i units with buyer i and the rest at the spot are given by $p_i q_i + \int_0^{Q-q_i} p_S(\tilde{q}) d\tilde{q} - \kappa Q$, the second derivative of seller profits with respect to quantity traded with buyer i is $p'_S(Q - q_i) < 0$.[9] Therefore, the seller is risk averse with respect to the quantity traded with buyer i. If the quantity contract instead specified that the expected trade is Y_i, but the variance of trade is $\gamma(1-\gamma)Y_i^2$, then the seller's expected profits would be smaller than $\mathbb{E}[\Pi^C]$.

However, if there is variation in market conditions, quantity contracts come at a cost:

- **Inflexibility.** The third and fourth panels of Figure 6 introduce the costs of inflexibility. Suppose that the intercept of the spot price curve $p_S(\tilde{q})$ is a random variable that captures market conditions. The third panel plots one curve where this intercept is unusually large. If the seller does not sign a contract with buyer i, then it never chooses to trade with buyer i upon observing this spot price curve. In contrast, if the seller signs a contract with buyer i, then it is required to trade Y_i units. The fourth panel shows that the seller would ex-post regret signing the contract, highlighting the costs of inflexibility.

We do not say much about buyer profits, either here or in our empirical analysis. If buyers also get dispersed price quotes at the spot market and hence face an upward-sloping spot price curve, they could be similarly risk-averse relative to the quantities traded in contracts. These considerations only potentially increase the value of quantity contracts.[10]

The two key assumptions are critical. Absent spot price dispersion, there would be no sense in which some buyers are more desirable than others, shutting down the buyer selection, quantity insurance, and inflexibility channels. Furthermore, if production was elastic, then the seller would trade with all spot buyers that have willingness to pay greater than marginal production cost κ, regardless of how much it trades with buyer i. This shuts down the quantity insurance and inflexibility channels.

5 Empirical Model

In this section, we develop an empirical model of spot trade and contract trade. The purpose of the model is to establish structure for estimation and counterfactual simulation.

In summary, we model trade as a game between a large seller and many smaller buyers. The buyers are either contract buyers or spot buyers, depending on whether they trade through quantity contracts. Importantly, we assume the seller's total sales to contract and spot buyers is exogenous in the short-run. The game proceeds according to the following timeline:

[9] We do not assume differentiability in the full model.

[10] We do not observe buyer-level data and hence do not estimate the effect of quantity contracts on buyers' profits.

1. At the start of the year, the seller negotiates the terms of quantity contracts with the contract buyers. Due to contracting frictions, these contracts are not fully contingent on the realizations of buyer and seller shocks. The contracted quantities are set to maximize joint pairwise expected surplus. Rebates are bargained over using Nash-in-Nash bargaining. Outside options for buyers and sellers are given by their expectations over future monthly spot markets.

2. Each month, random variables including production, market price, and buyers' outside options realize.

3. The seller trades with contract buyers as specified in the quantity contracts.

4. The seller trades all of its remaining production with spot buyers. The spot market is decentralized and due to trading frictions the seller meets only a fraction of all possible buyers. The seller makes take-it-or-leave-it offers to a subset of the buyers it meets. The offer depends on the exogeneously given outside option of the buyer. Variation in these outside options generates price dispersion in spot market.

In the remainder of the section, we elaborate on the details of the game. We start by specifying the environment: buyers have heterogeneous outside options and the seller has production that is inelastic in the short-run. We then develop the spot trade game and explain how expectations about the monthly spot market affect the quantities and rebates in annual quantity contracts. Finally, we provide some additional empirical specifications such as distributional assumptions.

5.1 Environment

There is a large seller and many smaller buyers. The seller and buyers operate in a few markets defined by the combination of global region (Region A or Region B) and fiber (hardwood or softwood). The seller has production Q_{jt} in market j and month t. We assume that the seller cannot re-allocate its production across regions within a month. Buyer i in market j and month t has demand Y_{ijt}. We refer to Y_{ijt} as the buyer's 'quantity type.'

There are three main outcome variables: a trade indicator, rebate, and quantity. The variable $\tau_{ijt} \in \{0,1\}$ indicates whether buyer i in market j and month t trades with the seller. If trade occurs, the buyer receives rebate R_{ijt} and purchases quantity q_{ijt}. The price measure is the mill gate price p_{ijt}, which depends on the rebate R_{ijt}, the market price $\overline{p}_{jt}$, and logistics costs c_{ijt} as follows: $p_{ijt} = \overline{p}_{jt}(1 - R_{ijt}) - c_{ijt}$.

Our first key assumption is that there is dispersion in the buyers' outside options. If trade does not occur between a buyer i and the seller in market j and month t, then the buyer can purchase its demand Y_{ijt} at rebate θ_{ijt}. More generally, if the buyer trades $q_{ijt} \in [0, Y_{ijt}]$ with the seller, then the buyer can purchase its residual demand $Y_{ijt} - q_{ijt}$ at rebate θ_{ijt}. To support this assumption, Figure 1 documents substantial price dispersion among spot buyers, consistent with dispersion in buyer outside options. We refer to θ_{ijt} as the 'rebate type' of buyer i in market j and month t.

Our second key assumption is that production and demand are inelastic at the monthly level. The seller and the buyers cannot adjust production Q_{jt} and demand Y_{ijt} in response to market conditions in the short term. The elasticity estimates in Table 2 support this assumption. We abstract from inventory and assume that the seller sells (at most) its entire production and the buyer purchases (at most) its entire demand.

The seller's payoff in market j and month t is its profit, assuming constant marginal production cost κ_j:

$$\Pi_{jt} = \sum_i \tau_{ijt}(\overline{p}_{jt}(1 - R_{ijt}) - c_{ijt})q_{ijt} - \kappa_j Q_{jt}.$$

Buyer i's payoff in market j and month t is its profit, assuming constant marginal valuation v_{ij}:[11]

$$U_{ijt} = \begin{cases} v_{ij}Y_{ijt} - \overline{p}_{jt}(1 - R_{ijt})q_{ijt} - \overline{p}_{jt}(1 - \theta_{ijt})(Y_{ijt} - q_{ijt}) & \text{if } \tau_{ijt} = 1, \\ v_{ij}Y_{ijt} - \overline{p}_{jt}(1 - \theta_{ijt})Y_{ijt} & \text{if } \tau_{ijt} = 0. \end{cases}$$

Variation in buyer rebate type θ_{ijt} may be driven by variation in valuation v_{ij} or by variation in bargaining power with other sellers. Since we observe only data from one seller, we do not separate these two sources of variation.

In each year y, buyers in each market j are partitioned into contract buyers in C_{jy} and spot buyers S_{jy}. In month t, the seller trades quantity Q^C_{jt} to contract buyers in market j. The seller trades the remaining $Q^S_{jt} = Q_{jt} - Q^C_{jt}$ tons with the spot buyers in market j. The determination of the three main outcome variables (trade indicator τ_{ijt}, rebate R_{ijt}, and quantity q_{ijt}) depends on whether a buyer is a spot buyer or contract buyer. We consider each case in turn.

5.2 Spot Trade

To summarize the spot game, the seller comes into contact with a fraction γ of the available spot buyers, where γ indexes the extent of trading frictions. Then, the seller trades with the most desirable spot buyers it comes into contact with until it runs out of production.

Due to trading frictions, the seller only comes into contact with a subset N_{jt} of the spot buyers in S_{jy} each month. The seller can only trade with spot buyers in N_{jt}. We assume that the set N_{jt} is selected

[11] We do not require marginal production cost κ_j or marginal valuation v_{ij} for counterfactual simulation. As such, we do not estimate these parameters.

arbitrarily from the set of spot buyers S_{jy}[12] and the cardinality is given by:

$$|N_{jt}| = \gamma |S_{jy}|.$$

That is, the seller comes in contact with a fraction γ of all spot buyers in the market. The fraction γ indexes the extent of trading frictions. As γ approaches one, trading frictions become minimal meaning that the seller can trade with all spot buyers in S_{jy}. Given that we construct S_{jy} as the set of spot buyers we observe in the data, we do not interpret a value of γ equal to one as the absence of trading frictions. There may be buyers outside of S_{jy} with whom the seller does not contact due to trading frictions even when γ equals one.

After the seller comes into contact with the spot buyers $N_{jt} \subseteq S_{jy}$, it makes take-it-or-leave-it rebate and quantity offers to the buyers in N_{jt}. Because production costs are sunk, we assume that the seller makes these offers to maximize total revenue from buyers in N_{jt}, subject to the constraints that total quantity sold to spot buyers in market j and month t is at most Q^S_{jt} and each buyer trades at most its demand Y_{ijt}:

$$\max_{\substack{\{q_{ijt}\} \\ \{r_{ijt}\}}} \sum_{i \in N_{jt}} (\overline{p}_{jt}(1 - R_{ijt}) - c_{ijt}) q_{ijt} \text{ such that } \sum_{i \in N_{jt}} q_{ijt} \leq Q^S_{jt} \text{ and } q_{ijt} \leq Y_{ijt} \text{ for all } i \in N_{jt}. \tag{1}$$

We assume that buyers choose whether to accept the offer in order to maximize payoff U_{ijt}.

In equilibrium, the seller trades as much as possible with the most desirable buyers in N_{jt}. The seller orders the buyers in N_{jt} in terms of mill gate price $\overline{p}_{jt}(1 - \theta_{ijt}) - c_{ijt}$. The seller offers a trade of quantity q_{ijt} equal to demand Y_{ijt} and rebate R_{ijt} equal to the rebate type θ_{ijt} for the buyers at the top of the list until the total quantity offered equals Q^S_{jt} (abstracting from any slack that arises due to discreteness). All buyers accept the offer.

The seller's expected profit from spot trade in this market $\Pi^S_{jt}(S_{jy}, Q^S_{jt})$ depends on the set of spot buyers S_{jy} and the total spot quantity Q^S_{jt}. The buyer's expected profit from spot trade in this market is given by U^S_{ij}.

5.3 Contract Trade

The model of contract trade follows the structure of quantity contracts in the industry and is motivated by the stylized facts presented in Section 3. At the start of the year, the seller and a contract buyer commit to always trade the buyer's demand. Furthermore, mill gate price is determined according to a fixed annual

[12] There are many search processes that could generate N_{jt}, such as directed search and sequential search. Because we do not observe the set of spot buyers with whom the seller comes into contact, we are unable to separately identify these processes. A limitation of our analysis is that contract buyers might come into contact with the seller with high probability when placed in the spot market, decreasing the value of contracting.

rebate. The buyer and seller bargain over the rebate via Nash-in-Nash bargaining.

A quantity contract is a map from the state space to the outcome space. The state is the realization of rebate type θ_{ijt} and quantity type Y_{ijt}. The outcome is the realization of the trade indicator τ_{ijt}, rebate R_{ijt}, and quantity q_{ijt}. We restrict the set of feasible contracts in a few ways:

1. First, we assume that trade always occurs, meaning that τ_{ijt} equals one. We relax this assumption when we fully specify the empirical model below.

2. Second, we specify that the rebate R_{ijt} is constant within a year y for a buyer i and market j: $R_{ijt} = R_{ijy}$. To justify this assumption, note that most annual contracts explicitly index price to the market price with a fixed rebate.

3. Third, we assume that quantity q_{ijt} equals the quantity type Y_{ijt}. Economic theory suggests that long-term relationships of many types may *increase* the magnitude of feasible quantity traded.[13] Empirically, the average contract buyer quantity is over five times the average spot buyer quantity. Therefore, assuming that quantity equals the quantity type and is not larger is a conservative assumption.

We assume that quantity contracts are perfectly enforceable. In practice, contracts often include 'contract fulfillment bonuses' that specify an additional rebate the buyer receives if it satisfies the terms of the contract. The data provider was unable to provide an example of a buyer that failed to satisfy the terms of the contract and lost the contract fulfillment bonus. Dynamics provide another enforcement mechanism. Many of the buyers have decades-long relationships with the seller, and risk losing a valuable trading partner if they renege. One could interpret the value of a quantity contract as the minimum amount that a buyer and seller would require to deviate from the terms of a contract, supposing the buyer and seller are playing a trigger strategy and have symmetric information. A positive value of contracts is sufficient for contracts to be self-enforceable under a symmetric information trigger strategy. We do not explicitly model these dynamic considerations for simplicity.

Due to the restrictions on the set of feasible quantity contracts, a contract can be fully described by the annual rebate R_{ijy}. The seller's expected profit from contract buyers in market j and month t depends on the set of rebates R_{ijy} for contract buyers in C_{jy} and equals:

$$\Pi^C_{jt}(\{R_{ijy}\}) = \sum_{i \in C_{jy}} \mathbb{E}[(\overline{p}_{jt}(1 - R_{ijy}) - c_{ijt})Y_{ijt} - \kappa_j Q^C_{jt}].$$

The expectation is over the quantity type Y_{ijt}. The buyer's expected profit depends on the rebate R_{ijy} and

[13] Macchiavello and Morjaria [2015] quantify this effect in the context of the Kenyan rose market.

equals:

$$U_{ij}^{C}(R_{ijy}) = \mathbb{E}[v_{ij}Y_{ijt} - (\overline{p}_{jt}(1-R_{ijy}) - c_{ijt})Y_{ijt}].$$

The expectation is over the quantity type Y_{ijt}.

We assume that a contract buyer and the seller choose the rebate R_{ijy} through a Nash-in-Nash bargaining game at the beginning of the year. Buyers and the seller have symmetric information about the distribution of types and market conditions. If bargaining fails, then the buyer joins the pool of spot buyers S_{jy}. As the equilibrium of Nash-in-Nash rebate negotiations, contract buyers and the seller choose the rebate R_{ijy} to solve:

$$\max_{R_{ijy}} \left(U_{ij}^{C}(R_{ijy}) - U_{ij}^{S}\right)^{\delta_{ijy}} \times \tag{2}$$

$$\left(\mathbb{E}\left[(\overline{p}_{jt}(1-R_{ijy}) - c_{ijt})Y_{ijt} + \Pi_{jt}^{S}\left(S_{jy}, Q_{jt} - \sum_{i_0 \in S_{jy}} Y_{i_0jt}\right) - \Pi_{jt}^{S}\left(S_{jy} \cup i, Q_{jt} - \sum_{i_0 \in S_{jy} \setminus i} Y_{i_0jt}\right)\right]\right)^{1-\delta_{ijy}}$$

such that $U_{ij}^{C}(R_{ijy}) \geq U_{ij}^{S}$ and

$$\mathbb{E}\left[(\overline{p}_{jt}(1-R_{ijy}) - c_{ijt})Y_{ijt} + \Pi_{jt}^{S}\left(S_{jy}, Q_{jt} - \sum_{i_0 \in S_{jy}} Y_{i_0jt}\right)\right] \geq \mathbb{E}\left[\Pi_{jt}^{S}\left(S_{jy} \cup i, Q_{jt} - \sum_{i_0 \in S_{jy} \setminus i} Y_{i_0jt}\right)\right].$$

The parameter δ_{ijy} is the relative bargaining power of buyer i in market j and year y. The first factor of (2) is the value of contracting for the buyer, given that the rebate is R_{ijy}. This factor equals the buyer's expected profit if it signs the contract minus the buyer's expected profit if it trades on the spot. The second factor of (2) is the value of contracting for the seller, given that the rebate is R_{ijy}. The first term of this factor equals the seller's expected revenue from the buyer if it signs the contract. The second term equals the seller's expected spot revenue if the buyer signs the contract. The third term equals the seller's expected spot revenue if the buyer rejects the contract and instead trades on the spot market. The maximization problem also requires that the ex-ante value of contracting is non-negative for both the buyer and seller.

If the constraints are non-binding, the solution to (2) is given by:

$$\begin{aligned}\mathbb{E}[(\overline{p}_{jt}(1-R_{ijy}) - c_{ijt})Y_{ijt}] = {} & (1-\delta_{ijy})\mathbb{E}[(\overline{p}_{jt}(1-\theta_{ijt}) - c_{ijt})Y_{ijt}] \\ & + \delta_{ijy}\left(\mathbb{E}\left[\Pi_{jt}^{S}\left(S_{jy} \cup i, Q_{jt} - \sum_{i_0 \in S_{jy} \setminus i} Y_{i_0jt}\right) - \Pi_{jt}^{S}\left(S_{jy}, Q_{jt} - \sum_{i_0 \in S_{jy}} Y_{i_0jt}\right)\right]\right).\end{aligned} \tag{3}$$

If the buyer has no relative bargaining power and δ_{ijy} equals zero, then the seller will set the rebate R_{ijy} such that the buyer's expected payoff equals its expected payoff if they are a spot buyer. This is also the expected payoff if the buyer never trades with the seller because the seller perfectly price discriminates with the spot buyers. If the rebate type θ_{ijt} and the quantity type Y_{ijt} are independent and the buyer has no bargaining power, then the seller will set the rebate R_{ijy} equal to the expected rebate type $\mathbb{E}[\theta_{ijt}]$. When the buyer's relative bargaining power δ_{ijy} is positive, the rebate R_{ijy} will be larger than the expected rebate type $\mathbb{E}[\theta_{ijt}]$ for contract buyers.

5.4 Empirical Specification

In this section, we describe two additional components of the model that facilitate estimation. First, we introduce an additional dimension of buyer type that explains why contract buyers occasionally do not trade. Second, we specify distributional assumptions on the buyer types.

The full empirical model includes an additional dimension of buyer type, the trade type ϕ_{ijt} for buyer i in market j and month t. We assume that, for contract buyers, the contract specifies that trade occurs if and only if the trade type ϕ_{ijt} is positive. The trade type does not enter into the spot game. The purpose of the trade type is to explain the empirical fact that the average contract buyer does not trade in 11% of months. If we were to instead assume that contract buyers trade with probability one, then we would overestimate the value of quantity contracts.[14]

Second, we impose additional structure on the buyer type distribution. We first decompose the three elements of buyer type into components that depends linearly on covariates and residuals:

$$\begin{aligned}
\theta_{ijt} &= X_{ijt}\beta_\theta + \epsilon^\theta_{ijt}, \\
Y_{ijt} &= C_{ijt}X^C_{ijt}\beta^C_Y + (1 - C_{ijt})X_{ijt}\beta^S_Y + \epsilon^Y_{ijt}, \\
\phi_{ijt} &= X_{ijt}\beta_\phi + \epsilon^\phi_{ijt}.
\end{aligned}$$

In this specification, the main covariate vector X_{ijt} includes a priority indicator, logistics costs, market price, a time trend in months, and market fixed effects. The secondary covariate vector X^C_{ijt} includes logistics costs, market price, a time trend in months, and a buyer-market-year fixed effect. The variable C_{ijt} indicates whether a buyer i is a contract buyer in market j and month t.

In summary, the rebate type θ_{ijt} and the trade type ϕ_{ijt} depend linearly on the main covariate vector. For spot buyers, the quantity type Y_{ijt} also depends linearly on the main covariate vector. For contract buyers, the quantity type Y_{ijt} instead depends on the secondary covariate vector X^C_{ijt}. Contract buyers tend to trade much larger quantities than spot buyers, so we include this additional heterogeneity into the

[14]In practice, contract buyers may sometimes not trade due to mill maintenance or inventory management. Contract-fulfillment rebates are paid based on total annual quantity, but there is no formal mechanism to enforce stable monthly quantities. Our model abstracts from such considerations.

quantity type. Furthermore, contract buyers tend to trade at vastly different scales, so it is important to incorporate buyer-market-year fixed effects in the quantity type.

We impose distributional assumptions on the residuals ϵ^{θ}_{ijt}, ϵ^{Y}_{ijt}, and ϵ^{ϕ}_{ijt}. We assume that the rebate type residual ϵ^{θ}_{ijt} and the quantity type residual ϵ^{Y}_{ijt} are jointly normal and i.i.d. distributed (we suppress conditioning for notational convenience):

$$\begin{pmatrix} \epsilon^{\theta}_{ijt} \\ \epsilon^{Y}_{ijt} \end{pmatrix} \sim \mathcal{N}\left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} \sigma_{\theta\theta} & \sigma_{\theta Y} \\ \sigma_{\theta Y} & C_{ijt}\sigma^{C}_{YY} + (1 - C_{ijt})\sigma^{S}_{YY} \end{pmatrix} \right).$$

Just as we incorporate additional heterogeneity in the mean of the quantity type by contract status, we also allow for heterogeneity in the residual variance of the quantity type by contract status. It is important that we correctly capture the variance of contract buyer quantity to quantify the value of quantity insurance.

We assume that the trade type residual ϵ^{ϕ}_{ijt} is i.i.d. type I extreme value (we suppress conditioning for notational convenience):

$$\epsilon^{\phi}_{ijt} \sim TIEV.$$

We also assume that the trade type residual ϵ^{ϕ}_{ijt} is independent of the rebate type residual ϵ^{θ}_{ijt} and the quantity type residual ϵ^{Y}_{ijt}. As justification, we note that Table OA.1 shows that the probability of trade is statistically flat over the average mill gate price for contract buyers.

6 Estimation

In this section, we describe the estimation routine. We explain how the parameters of the model are identified. We then present the main estimates.

We estimate ten parameter vectors: $(\beta_\theta, \beta^{C}_{Y}, \beta^{S}_{Y}, \beta_\phi, \sigma_{\theta\theta}, \sigma^{C}_{YY}, \sigma^{S}_{YY}, \sigma_{Y\theta}, \gamma, \delta_{ijy})$. The first eight parameter vectors govern the joint type distribution. The ninth parameter γ indexes the extent of trading frictions. The tenth parameter vector δ_{ijy} is the relative bargaining power of buyer i in market j and year y. We allow for heterogeneity in δ_{ijy} by quartiles of expected quantity type for contract buyers $X^{C}_{ijt}\beta^{C}_{Y}$.

We estimate the parameters in two steps. We provide further details in Online Appendix B.

Step 1: Spot Buyer Sample. The first step uses the spot buyer data to estimate the parameters that govern the rebate and quantity type distribution for spot buyers as well as the extent of trading frictions: $(\beta_\theta, \beta^{S}_{Y}, \sigma_{\theta\theta}, \sigma^{S}_{YY}, \sigma_{Y\theta}, \gamma)$. To estimate these parameters, we use two-step generalized method of moments. We match the distribution of rebates, quantity, and trade indicator among spot buyers to the model-implied values conditional on X_{ijt}. We have six moment vectors that all equal zero in expectation:

$$
\begin{aligned}
&\left(R_{ijt} - \mathbb{E}[\theta_{ijt}|\tau_{ijt} = 1, X_{ijt}]\right) X_{ijt}\tau_{ijt} \\
&\left(q_{ijt} - \mathbb{E}[Y_{ijt}|\tau_{ijt} = 1, X_{ijt}]\right) X_{ijt}\tau_{ijt} \\
&\left((R_{ijt} - \mathbb{E}[\theta_{ijt}|\tau_{ijt} = 1, X_{ijt}]])^2 - \mathrm{Var}(\theta_{ijt}|\tau_{ijt} = 1, X_{ijt}])\right) \tau_{ijt} \\
&\left((q_{ijt} - \mathbb{E}[Y_{ijt}|\tau_{ijt} = 1, X_{ijt}]])^2 - \mathrm{Var}(Y_{ijt}|\tau_{ijt} = 1, X_{ijt}])\right) \tau_{ijt} \\
&\left(R_{ijt}q_{ijt} - \mathbb{E}[\theta_{ijt}Y_{ijt}|\tau_{ijt} = 1, X_{ijt}]\right) \tau_{ijt} \\
&(\tau_{ijt} - \mathbb{E}[\tau_{ijt}|X_{ijt}])X_{ijt}
\end{aligned}
$$

The first five moment vectors match the mean and variance-covariance matrix of rebate and quantity to the model-implied values conditional on X_{ijt}. The sixth moment vector matches trade patterns to the model-implied value conditional on X_{ijt}. To calculate the model-implied values, we develop an approximation that we simulate using quadrature rules.

Step 2: Contract Buyer Sample. The second step uses the contract buyer data to estimate the remaining parameters: $(\beta_Y^C, \beta_\phi, \sigma_{YY}^C, \delta_{ijy})$. To estimate these parameters, we run linear and logistics regressions. We condition on the parameters estimated in the first step. We estimate the distribution of the trade type ϕ_{ijt} using logistic regression of the trade indicator on the main covariates X_{ijt} for contract buyers. We estimate the distribution of quantity type Y_{ijt} for contract buyers using ordinary least squares on the set of observations where trade occurs. The only remaining parameter is the relative bargaining power δ_{ijy}. To estimate δ_{ijy}, we first simulate the terms in (3) using the parameters that have already been estimated. We then estimate δ_{ijy} using ordinary least squares based on equation (3).

6.1 Identification

Table 4 summarizes the variation in the data that identifies the parameters of the model. To identify the distribution of the rebate type θ_{ijt}, we use the covariance between rebate conditional on trade and the main covariate vector X_{ijt}. For example, the first column shows that prioritized buyers have significantly lower rebates than other buyers. As a result, these buyers also have lower rebate type. The first column also shows that contract buyers have a significantly higher rebate than other buyers conditional on X_{ijt}. This difference identifies the relative bargaining power δ_{ijy}. The larger this increase, the larger δ_{ijy}.

The correlation between trade and the main covariate vector X_{ijt} for spot buyers identifies the extent of trading frictions γ. If trading frictions are large and the seller comes into contact with relatively few spot buyers, then the seller has to trade with every buyer with whom it comes into contact. Since the set of buyers

Table 4: Identifying Variation

	(1) Rebate	(2) Trade	(3) Trade	(4) Quantity	(5) Quantity
Prioritized	-0.04	0.15	0.01	-0.41	
	(0.002)	(0.01)	(0.01)	(0.03)	
Logistics Costs	-0.05	-0.03	-0.09	-0.33	1.49
	(0.005)	(0.03)	(0.03)	(0.07)	(0.56)
Market Price	0.03	0.01	0.02	0.05	-0.10
	(0.0008)	(0.006)	(0.004)	(0.02)	(0.03)
Time Trend (months)	-0.002	-0.0005	-0.001	0.0005	0.006
	(0.00005)	(0.0004)	(0.0003)	(0.001)	(0.007)
Contract Buyer	0.04				
	(0.002)				
Observations	6,492	5,964	3,600	3,256	3,236
R-squared	0.60	0.06	0.07	0.08	0.88
Sample	Full	Spot	Contract	Spot	Contract
Market FE	YES	YES	YES	YES	NO
Buyer-market-year FE	NO	NO	NO	NO	YES

Notes. Robust standard errors in parentheses. Rebate is defined as one minus price after rebates over market price. Quantity is in thousands of tons. Logistics costs is in hundreds of euros. Market price is in hundreds of euros. A market is defined as a product (hardwood or softwood) and region (Region A or Region B).

that the seller comes into contact with is chosen uniformly at random, there would be no correlation between trade and the main covariate vector X_{ijt}. However, the second column of Table 4 shows that prioritized spot buyers are significantly more likely to trade. Furthermore, Figure 2 shows that the probability of trade increases with average price for spot buyers. These results suggest that the seller can be somewhat selective about the spot buyers with whom it trades. Therefore, γ is larger than the minimum possible value.

Next, the correlation between trade and the main covariate vector X_{ijt} for contract buyers identifies the distribution of trade type ϕ_{ijt}. The correlation between quantity and the main covariate vector X_{ijt} identifies the distribution of the quantity type Y_{ijt} for spot buyers. Finally, the correlation between quantity and the secondary covariate vector X_{ijt}^C identifies the distribution of the quantity type Y_{ijt} for contract buyers.

6.2 Results

Table 5 presents the estimates of the model parameters. The parameters that govern the distribution of rebate type, quantity type, and trade type correspond in sign to the identifying variation in Table 4. For example, the estimates of β_θ corresponds to the coefficients in the first column of Table 4. The magnitudes differ from the regression version because the realized spot rebates are a certain truncation of the the rebate type distribution. Furthermore, the estimate of the distribution of quantity type for contract buyers is exactly the fifth column of Table 4. This is true because we estimate the quantity type for contract buyers with ordinary least squares.

Table 5: Estimates

Parameter	Rebate type θ_{ijt} β_θ	Quantity type Y_{ijt} (spot) β_Y^S	Quantity type Y_{ijt} (contract) β_Y^C	Trade type ϕ_{ijt} β_ϕ
Prioritized	-0.09	-0.39		0.13
Logistics Costs	-0.06	-0.39	1.49	-1.17
Market Price	0.02	0.50	-0.10	0.15
Time Trend (months)	-0.001	0.001	0.006	-0.01
Market FE	YES	YES	NO	YES
Buyer-market-year FE	NO	NO	YES	NO

Trading frictions γ	0.78
Buyer bargaining power δ_{ijy}: 1st quartile quantity type	0.22
Buyer bargaining power δ_{ijy}: 2nd quartile quantity type	0.77
Buyer bargaining power δ_{ijy}: 3rd quartile quantity type	0.34
Buyer bargaining power δ_{ijy}: 4th quartile quantity type	0.17

Notes. Quantity type is in thousands of tons. Logistics costs is in hundreds of euros. Market price is in hundreds of euros. A market is defined as a product (hardwood or softwood) and region (Region A or Region B).

We estimate that the extent of trading frictions γ equals 0.78. This means that the seller comes into contact with 78% of the spot buyers in each month. For comparison, the average probability of trade among spot buyers is 65%. Therefore, the seller is somewhat selective about the spot buyers with whom it trades. However, trading frictions are also not minimal because we estimate γ less than one.

Finally, we estimate that buyer bargaining power δ_{ijy} is 0.38 on average. Institutional details suggest that the seller should have more than half of the relative bargaining power because the seller is larger than most buyers. There is also substantial heterogeneity in buyer bargaining power depending on the quartile of quantity type. All the estimates are between zero and one. Note that the buyer bargaining power δ_{ijy} is the bargaining power when signing a quantity contract. There may be additional bargaining power that enters into the determination of the rebate type θ_{ijt}.

7 Counterfactual Analysis

In this section, we present the results of three counterfactual analyses. The first counterfactual quantifies the value of quantity contracts and documents comparative statics. The second counterfactual isolates the value of quantity contracts due to quantity insurance and the cost due to inflexibility. The third counterfactual quantifies how the use of quantity contracts changes with the extent of trading frictions.

Counterfactual 1. Value of Quantity Contracts and Comparative Statics. Figure 7 presents a histogram of the value of quantity contracts, where each observation is a contract buyer in a particular market and year. The value of a contract is defined as the difference between buyer and seller profits in the observed allocation and a counterfactual allocation where that buyer trades on the spot.[15] The median value of a quantity contract is €57/ton, and the value is positive for 87% of contracts. To put the magnitude in context, this equals 10% of the average price after rebates. There is also substantial heterogeneity in the value of contracts. Because we quantify the value of contracts in euros per tons traded, some of this heterogeneity is due to heterogeneity in the denominator, quantity traded. The value of contracts may not be multiplicative in the quantity traded.

Figure 7: Quantity Contracts Are Valuable

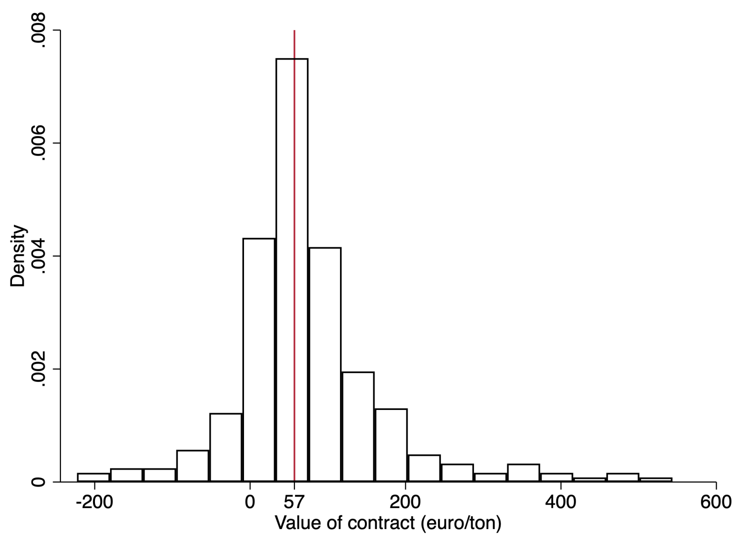

Notes. The median value is 57 €/ton. An observation is a contract buyer within a market and year. This figure includes the middle 96% of contracts.

To explore the determinants of the value of quantity contracts, we present four comparative statics in Figure 8. First, Figure 8a shows that the value of contracts decreases with contract buyer logistics costs. For each multiplier on the x-axis, we calculate the value of contracts after multiplying every contract buyer's logistics costs by that multiplier. This comparative static highlights how contracts get their value from the buyer selection channel. Contracts allow the seller to lock-in trade with the most desirable trading partners.

Second, Figure 8b shows that the value of quantity contracts decreases with the variance of contracted

[15] We simulate the spot outcome using an approximation to the seller's optimization problem (1). We specify that the seller orders the spot buyers it comes into contact with by mill gate price. The seller issues offers to trade according to this order until the next offer would lead to a violation of the seller's quantity constraint. This approximation may lead to excess slack in the seller's quantity constraint relative to the solution to (1). In particular, if a large spot buyer is the first spot buyer that the seller does not trade with (according to the ordering by mill gate price), then there could be significant slack. Most contract buyers are large, so the algorithm artificially depresses the seller's profits in a counterfactual where a contract buyer trades on the spot. Therefore, the approximation unduly increases the value of contracts and is a limitation of our analysis.

quantity. For each multiplier on the x-axis, we calculate the value of contracts after multiplying the residual component of contract quantity type by the square root of the multiplier. When variance of quantity is only 25% higher than the baseline, the median contract is not valuable. This result highlights how contracts get their value from the quantity insurance channel. The seller is endogenously risk averse with respect to quantity traded with a contract buyer.

Figure 8: Comparative Statics of the Value of Quantity Contracts

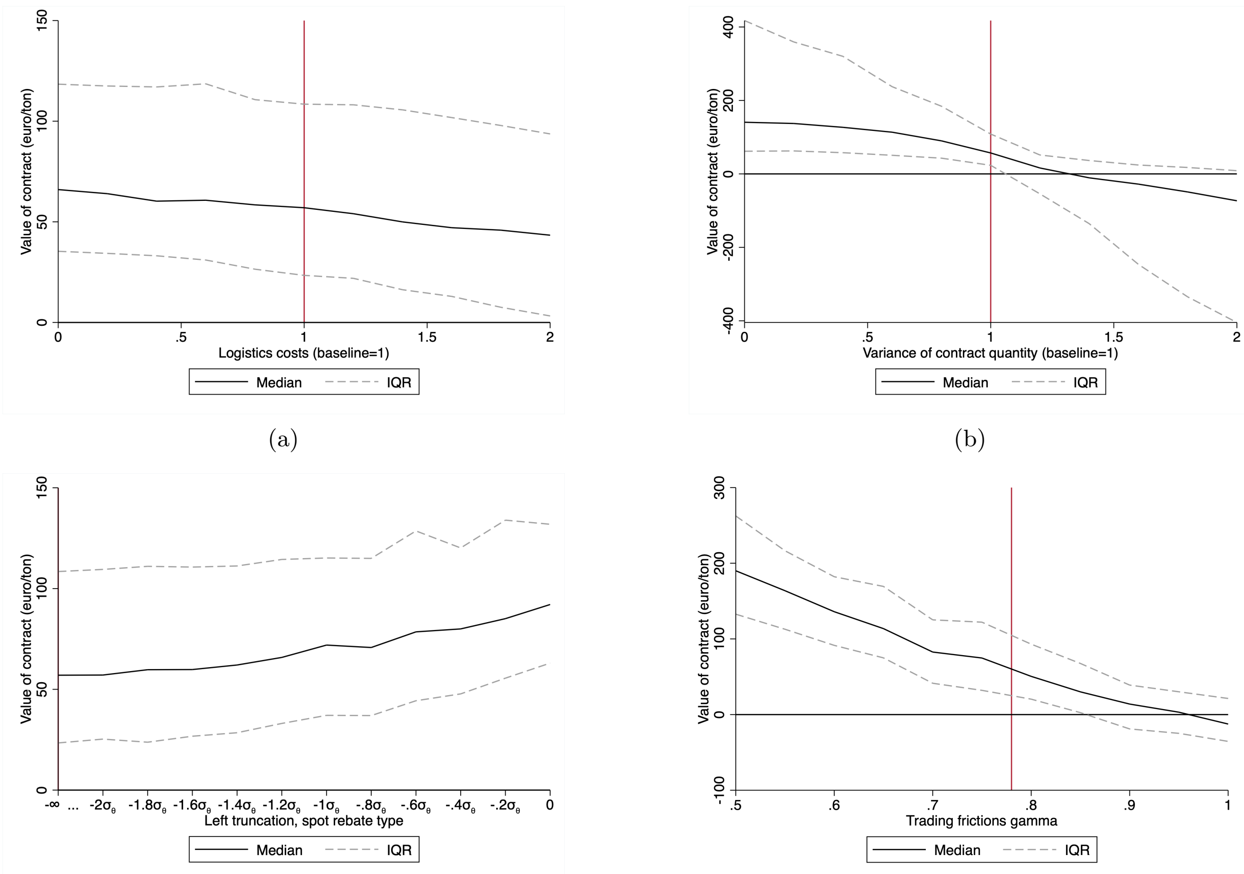

(a) (b) (c) (d)

Third, Figure 8c shows that the value of quantity contracts increases in the left truncation of the spot buyer rebate type distribution. The baseline is a truncation of $-\infty$. A left truncation on the spot buyer rebate type distribution decreases the probability that the seller would ex-post prefer to allocate more quantity to the spot and away from contract buyers. This result highlights the cost of inflexibility.

Fourth, Figure 8d shows that the value of quantity contracts decreases with the extent of trading frictions γ. In fact, when trading frictions are minimal, the median contract is not valuable. This result highlights how trading frictions are necessary for contracts to have value. When trading frictions are minimal and γ equals one, the seller can trade with any of the spot buyers, so it has maximal choice. In contrast, if the seller signs a contract, then it restricts its choice set, because it is forced to trade with the contract buyer. As a result, the seller prefers to trade on the spot.[16]

Counterfactual 2. Isolation of Quantity Insurance and Inflexibility. There are many ways to remove quantity insurance from the value of quantity contracts. The model predicts that the value of contracts decreases in the variance of contracted quantity. This comparative static motivates our chosen approach. We change the variance of a contract buyer's quantity to equal what it would be if that buyer instead traded on the spot. For 70% of contracts, the variance of quantity increases when they counterfactually trade on the spot.[17] Figure 9 shows that the distribution of the value of contracts is left-shifted, and the median value is €43/ton. We thus conclude that, without quantity insurance, quantity contracts would be 25% less valuable.

Figure 9: Quantity Contracts Are 25% Less Valuable Without Quantity Insurance

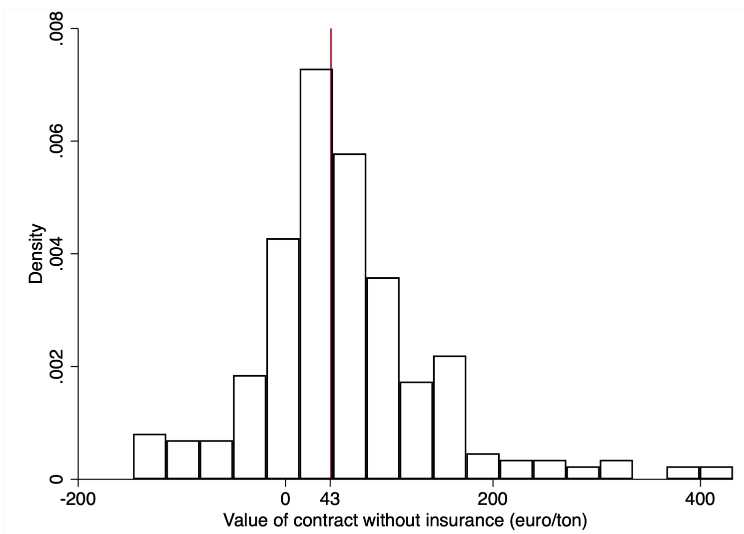

Notes. The median value is 43 €/ton. An observation is a contract buyer within a market and year. This figure includes the middle 96% of contracts.

Next, we remove the cost of inflexibility. To do so, we allow the seller to ex-post renege on its contracted quantity, guaranteeing full flexibility to market conditions.[18] Figure 10 presents the resulting distribution of the value of quantity contracts. All values are positive, consistent with the idea that quantity contracts only increase the seller's choice set absent the cost of inflexibility. Furthermore, the median value increases to €105/ton due to the increased option value. We conclude that, without the cost of inflexibility, quantity contracts would be 84% more valuable.

[16]The approximation to (1) causes some contracts have positive value even when trading frictions are minimal.

[17]One reason the variance may decrease for some contracts is that trade probabilities are so low counterfactually that quantity equals zero almost every month and hence the variance is close to zero.

[18]We approximate the result of optimization through an algorithm where the seller has the option to treat contract buyers either as an additional spot buyer or to follow the terms of the contract. This approximation weakly decreases the size of the seller's choice set, so it weakly decreases the value of quantity contracts without the cost of inflexibility. Therefore, we underestimate the cost of inflexibility.

Figure 10: Quantity Contracts Are 84% More Valuable Without Inflexibility

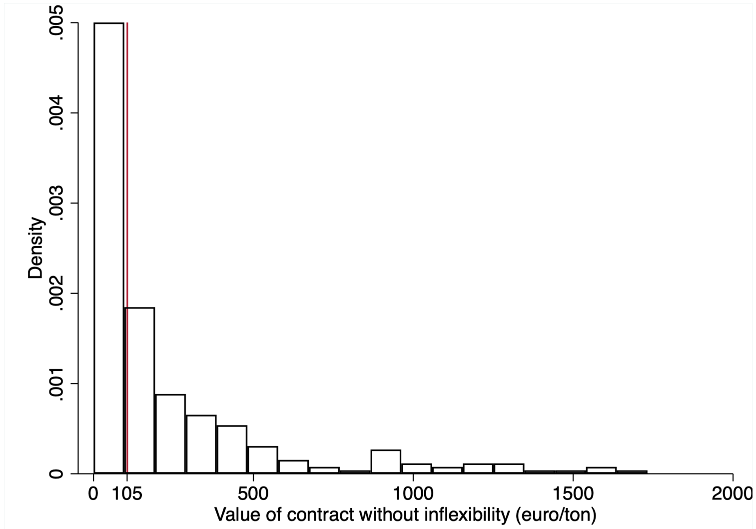

Notes. The median value is 105 €/ton. An observation is a contract buyer within a market and year. This figure includes the middle 96% of contracts.

Counterfactual 3. Contracting Behavior and Trading Frictions. The third counterfactual evaluates how contracting behavior changes with the extent of trading frictions. Intuitively, because quantity contracts are less valuable when trading frictions are low, we expect they would be used less. We have thus far refrained from specifying how buyers are allocated between contract and spot. For the purposes of this counterfactual, we assume that the seller chooses an allocation of buyers into contract and spot in order to maximize total profits. In practice, some buyers may refuse the seller's request to write a contract because they already have written large contracts with other sellers, and thus prefer the flexibility of trading on the spot. We abstract from these considerations in this analysis.[19]

Less quantity would be contracted as trading frictions diminish. At the estimated value of γ equal to 0.78, the allocation algorithm predicts that 84.9% of quantity would be contracted. In the data, 84.6% of quantity is contracted. However, the set of contract buyers that the algorithm predicts differs from the set of contract buyers we observe in the data. Figure 11 shows that as trading frictions diminish, the proportion of quantity contracted decreases. When γ is less than 0.6, almost all quantity is contracted, but only 30% of quantity is contracted when trading frictions are minimal and γ equals one.

Figure 12 shows that the seller's total profits increase as trading frictions diminish. The seller's profits are 10% higher if trading frictions are minimal. This profit is the predicted value when the seller can re-allocate

[19] A further complication is that it is too computationally demanding to calculate the seller's profits under all possible allocations of buyers into spot and contract, because the number of possible allocations is exponential in the number of buyers. Therefore, we suppose that the seller uses the following algorithm to assign buyers into spot and contract. The seller follows this algorithm in parallel for each market and year. First, start with all buyers in spot. Calculate total profits if exactly one spot buyer writes a contract. Allocate the spot buyer to contract that would imply the largest total profits. Next, return to the pool of remaining spot buyers. Repeat this algorithm until profits only decrease when signing an additional contract.

Figure 11: Less Quantity Is Contracted When Trading Frictions Diminish

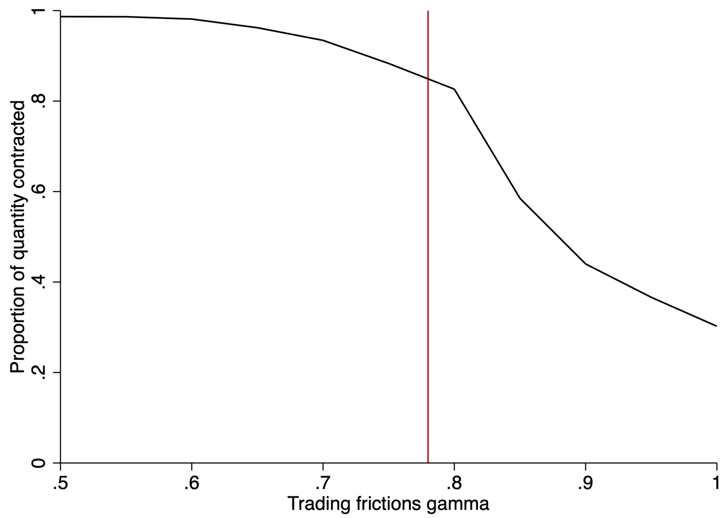

buyers at will between contract and spot. Furthermore, the proportion of that profits that comes from contract buyers decreases as trading frictions diminish. This figure suggests that efforts to reduce trading frictions would reduce the use of contracts and increase profits.

Figure 12: Profits Are Higher When Trading Frictions Diminish

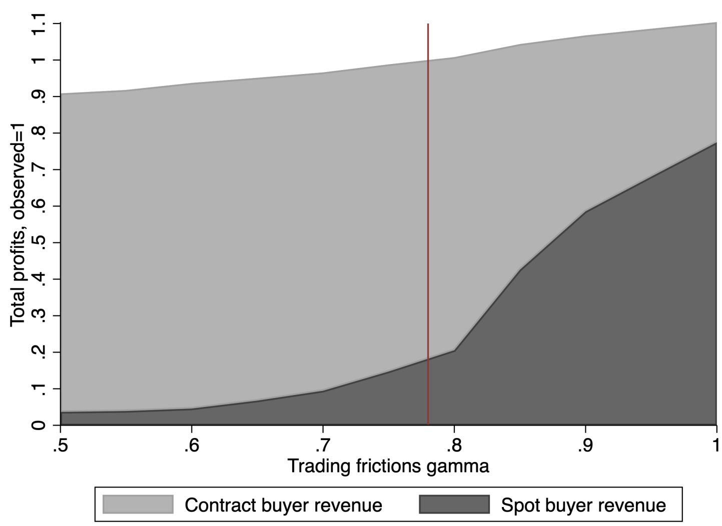

Finally, Figure 13 considers the seller's total profits over trading frictions when quantity contracts are not available. The main takeaway from the figure is that the slope is larger when quantity contracts are not available.[20] That is, the marginal effect of an increase in trading frictions on seller profits is larger when the seller is unable to use quantity contracts. Moving from the estimated trading frictions to minimal trading

[20]In principle, the ability to use quantity contracts should only increase seller profits by expanding the seller's choice set. However, in some cases we find that seller profits are higher without quantity contracts due to approximation for (1) we use.

frictions increases the seller's profits by 17% if it cannot use quantity contracts but only increases the seller's profits by 10% if it can use quantity contracts. Therefore, quantity contracts mitigate the costs of trading frictions.

Figure 13: Quantity Contracts Reduce the Costs of Trading Frictions

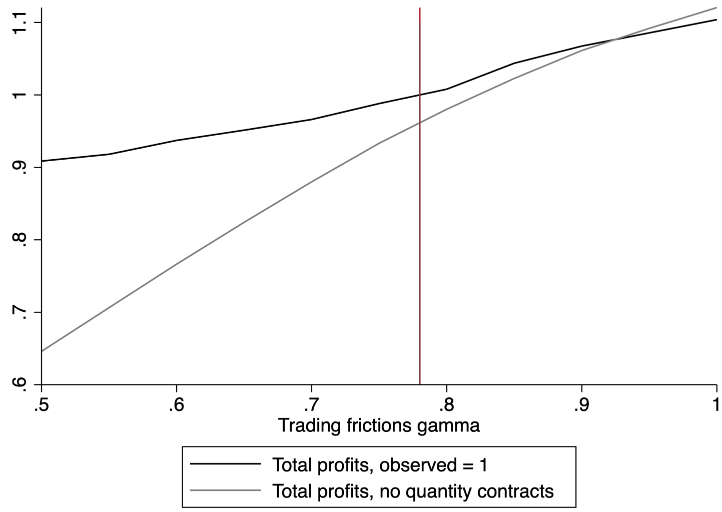

8 Discussion

Our research makes clear that there are multiple structures that buyers and sellers can use to organize trade in light of trading frictions. In this section, we discuss how certain markets might arrive at different structures. An interesting area of future research is to further develop a theory of how and why the equilibrium outcome arises.

Why do some industries rely on intermediaries and others rely on quantity contracts? First, industries where buyers and sellers are large and have bargaining power might rely on quantity contracts instead of intermediaries. Gavazza [2016] studies intermediation in the market for business jet aircraft, and one finding is that intermediaries do not enter if their relative bargaining power is low. Unlike in the market for business jet aircraft, buyers and sellers in the pulp industry tend to be quite large. Therefore, intermediaries may have low bargaining power and find it unprofitable to coordinate trade between large buyers and sellers. Consistent with this view, some small buyers in the pulp industry trade through intermediaries. Second, markets where buyers and sellers have highly capital-intensive production might rely on quantity contracts instead of intermediaries. Quantity contracts commit buyers and sellers to a quantity in advance, but intermediaries organize trade at the point of sale. Flexibility is less important for industries with capital-intensive production where total market participation is known well in advance. Consistent with this view, Salz [2020] finds that intermediation increases welfare in the market for trade-waste, where many buyers are firms that operate in the service sector. These buyers may require more flexibility because waste production is elastic to market conditions.

A related question is why industries use contracts of varying time-horizons. On one hand, shorter horizon contracts are more flexible to market conditions because they can be updated more frequently. On the other hand, longer horizon contracts may help prevent hold-up on long-term capital investment. Therefore, in industries with specific investment, such as the market for liquefied natural gas [Zahur, 2020], contracts may have longer horizon.

9 Conclusion

In many industries, buyers and sellers rely on quantity contracts to organize trade. Our main contribution is to show the value and cost of quantity contracts. First, quantity contracts increase trade between high surplus trading partners. Second, quantity contracts serve as a form of quantity insurance–quantity contracts protect buyers and sellers from the risk of having to trade with low surplus trading partners. However, quantity contracts are inflexible to changing market conditions. Our second contribution is to develop an empirical model of quantity contracts and quantify these forces in the pulp and paper industry. We find that the median value of a quantity contract is 10% of net price. Furthermore, most of this value comes from the buyer selection channel, though the quantity insurance channel is still sizable. The cost of inflexibility is large. If trading frictions diminish, quantity contracts are used less frequently and profits increase.

References

Philippe Aghion and Patrick Bolton. An incomplete contracts approach to financial contracting. *The review of economic Studies*, 59(3):473–494, 1992.

Patrick Bajari and Steven Tadelis. Incentives versus transaction costs: A theory of procurement contracts. *Rand journal of Economics*, pages 387–407, 2001.

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

Onur Boyabatlı, Paul R Kleindorfer, and Stephen R Koontz. Integrating long-term and short-term contracting in beef supply chains. *Management Science*, 57(10):1771–1787, 2011.

Yvonne Brodrechtova. Economic valuation of long-term timber contracts: Empirical evidence from Germany. *Forest Policy and Economics*, 55:1–9, 2015.

Ambarish Chandra and Mariano Tappata. Consumer search and dynamic price dispersion: an application to gasoline markets. *The RAND Journal of Economics*, 42(4):681–704, 2011.

Allan Collard-Wexler, Gautam Gowrisankaran, and Robin S Lee. “nash-in-nash” bargaining: A microfoundation for applied work. *Journal of Political Economy*, 127(1):163–195, 2019.

Kenneth S Corts and Jasjit Singh. The effect of repeated interaction on contract choice: Evidence from offshore drilling. *Journal of Law, Economics, and Organization*, 20(1):230–260, 2004.

Gregory S Crawford and Ali Yurukoglu. The welfare effects of bundling in multichannel television markets. *American Economic Review*, 102(2):643–85, 2012.

Mathias Dewatripont and Jean Tirole. A theory of debt and equity: Diversity of securities and manager-shareholder congruence. *The quarterly journal of economics*, 109(4):1027–1054, 1994.

Javier D Donna, Pedro Pereira, Tiago Pires, and Andre Trindade. Measuring the welfare of intermediaries in vertical markets. *SSRN 3285205*, 2019.

Mark Egan. Brokers versus retail investors: Conflicting interests and dominated products. *The Journal of Finance*, 74(3):1217–1260, 2019.

Maryam Farboodi, Gregor Jarosch, and Guido Menzio. Intermediation as rent extraction. Technical report, National Bureau of Economic Research, 2018.

Yuk-fai Fong and Jin Li. Relational contracts, limited liability, and employment dynamics. *Journal of Economic Theory*, 169:270–293, 2017.

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

Alessandro Gavazza. An empirical equilibrium model of a decentralized asset market. *Econometrica*, 84(5):1755–1798, 2016.

Alessandro Gavazza and Alessandro Lizzeri. Frictions and intermediation in product markets. *Prepared for the Handbook of Industrial Organization, vol. IV*, 2021.

Ricard Gil and Justin Marion. Self-enforcing agreements and relational contracting: Evidence from california highway procurement. *The Journal of Law, Economics, & Organization*, 29(2):239–277, 2013.

Ricard Gil and Giorgio Zanarone. On the determinants and consequences of informal contracting. *Journal of Economics & Management Strategy*, 27(4):726–741, 2018.

Matthew Grennan. Price discrimination and bargaining: Empirical evidence from medical devices. *American Economic Review*, 103(1):145–77, 2013.

Marina Halac. Relational contracts and the value of relationships. *American Economic Review*, 102(2):750–79, 2012.

Susan Helper and Rebecca Henderson. Management practices, relational contracts, and the decline of general motors. *Journal of Economic Perspectives*, 28(1):49–72, 2014.

Chang-Tai Hsieh and Enrico Moretti. Can free entry be inefficient? fixed commissions and social waste in the real estate industry. *Journal of Political Economy*, 111(5):1076–1122, 2003.

Paul L Joskow. Contract duration and relationship-specific investments: Empirical evidence from coal markets. *The American Economic Review*, pages 168–185, 1987.

Paul L Joskow. Price adjustment in long-term contracts: the case of coal. *The Journal of Law and Economics*, 31 (1):47–83, 1988.

Greg Kaplan and Guido Menzio. The morphology of price dispersion. *International Economic Review*, 56(4):1165–1206, 2015.

Christian Kimmich and Urs Fischbacher. Behavioral determinants of supply chain integration and coexistence. *Journal of Forest Economics*, 25:55–77, 2016.

Francine Lafontaine and Margaret Slade. Inter-firm contracts: Evidence, 2010.

Richard A Lambert. Long-term contracts and moral hazard. *The Bell Journal of Economics*, pages 441–452, 1983.

Phillip Leslie and Alan Sorensen. Resale and rent-seeking: An application to ticket markets. *Review of Economic Studies*, 81(1):266–300, 2014.

Jin Li and Niko Matouschek. Managing conflicts in relational contracts. *American Economic Review*, 103(6):2328–51, 2013.

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

James M Malcomson. Relational incentive contracts. In *The handbook of organizational economics*, pages 1014–1065. Princeton University Press, 2012.

Kurt Niquidet and Glen O'Kelly. Forest–mill integration: A transaction cost perspective. *Forest Policy and Economics*, 12(3):207–212, 2010.

Nancy Kane Ohanian. Vertical integration in the us pulp and paper industry, 1900-1940. *The Review of Economics and Statistics*, pages 202–207, 1994.

Martin Pesendorfer. Horizontal mergers in the paper industry. *RAND Journal of Economics*, pages 495–515, 2003.

Ariel Rubinstein and Asher Wolinsky. Equilibrium in a market with sequential bargaining. *Econometrica: Journal of the Econometric Society*, pages 1133–1150, 1985.

Tobias Salz. Intermediation and competition in search markets: An empirical case study. Technical report, National Bureau of Economic Research, 2020.

Victor Stango and Jonathan Zinman. Borrowing high versus borrowing higher: price dispersion and shopping behavior in the us credit card market. *The Review of Financial Studies*, 29(4):979–1006, 2016.

George J Stigler. The economics of information. *Journal of political economy*, 69(3):213–225, 1961.

Bruno Strulovici. Contract negotiation and the coase conjecture: A strategic foundation for renegotiation-proof contracts. *Econometrica*, 85(2):585–616, 2017.

Bart S Vanneste and Phanish Puranam. Repeated interactions and contractual detail: Identifying the learning effect. *Organization science*, 21(1):186–201, 2010.

Gewei Wang. *Does market concentration motivate pulp and paper mills to vertically integrate?* PhD thesis, Georgia Institute of Technology, 2005.

Oliver E Williamson. Transaction-cost economics: the governance of contractual relations. *The journal of Law and Economics*, 22(2):233–261, 1979.

Frank A Wolak. An empirical analysis of the impact of hedge contracts on bidding behavior in a competitive electricity market. *International Economic Journal*, 14(2):1–39, 2000.

Nahim Bin Zahur. Long-term contracts and efficiency in the liquefied natural gas industry. 2020.

Online Appendix

Pulp Friction: The Value of Quantity Contracts in Decentralized Markets

Olivier Darmouni Simon Essig Aberg Juha Tolvanen

Online Appendix A Additional Exhibits

Table OA.1: Patterns in Probability of Trade

	(1) Pr(trade)	(2) Pr(trade)	(3) Pr(trade)	(4) Pr(trade)
Contract Buyer	0.35	0.26	0.65	0.55
	(0.03)	(0.04)	(0.17)	(0.21)
Spot Buyer × Avg. Price			0.31	1.04
			(0.23)	(0.53)
Contract Buyer × Avg. Price			-0.28	0.51
			(0.25)	(0.60)
Capacity (Millions Tons)		0.12		0.15
		(0.38)		(0.38)
Prioritized		0.07		0.05
		(0.03)		(0.03)
Logistics Costs		-0.18		-0.13
		(0.09)		(0.11)
Observations	797	576	797	576
R-squared	0.13	0.19	0.13	0.20
Market-year FE	NO	YES	NO	YES

Notes. Robust standard errors in parentheses. Each observation is a buyer within a region, fiber, and year. Logistics costs are in hundreds of euros. Average price is mill gate price in 100,000 euros.

Figure OA.1: Example Spot Demand Curves

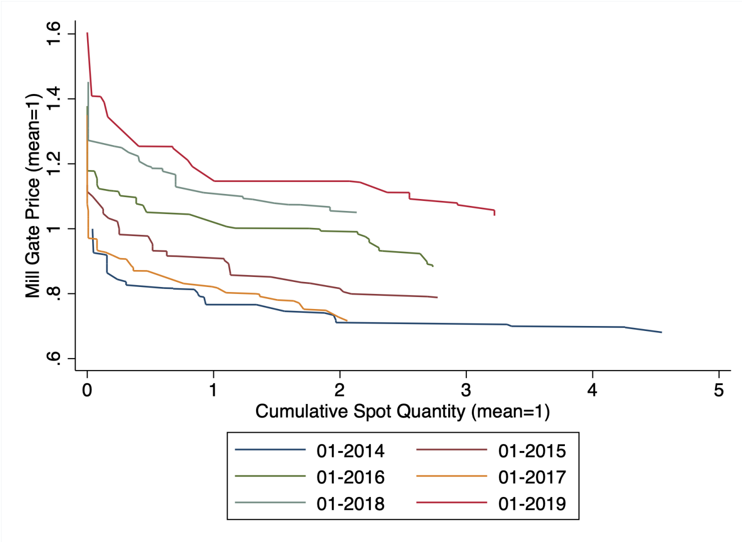

Notes. Each curve considers the distribution of mill gate prices and quantities within a particular region, fiber, and month. The lowest 5% of prices are excluded. The cumulative spot quantity equals cumulative quantity sold to the spot under the assumption that the first unit sold to the spot receives the highest price.

Table OA.2: Contract Buyers are More Reliable than Spot Buyers

	(1) C.V. quantity	(2) C.V. quantity	(3) C.V. quantity (cond. on trade)	(4) C.V. quantity (cond. on trade)
Contract buyer	-0.85	-0.69	-0.09	-0.09
	(0.05)	(0.06)	(0.02)	(0.02)
Capacity (millions tons)		-0.72		0.22
		(0.53)		(0.22)
Prioritized		-0.25		-0.03
		(0.06)		(0.02)
Observations	797	574	797	572
R-squared	0.25	0.37	0.03	0.10
Market-year FE	NO	YES	NO	YES

Notes. Robust standard errors in parentheses. 'C.V.' stands for the coefficient of variation, defined as the standard deviation divided by the mean. Each observation is a buyer within a region, fiber, and year.

Figure OA.2: Spot Buyers Are As Reliable as Contract Buyers Conditional on Trade

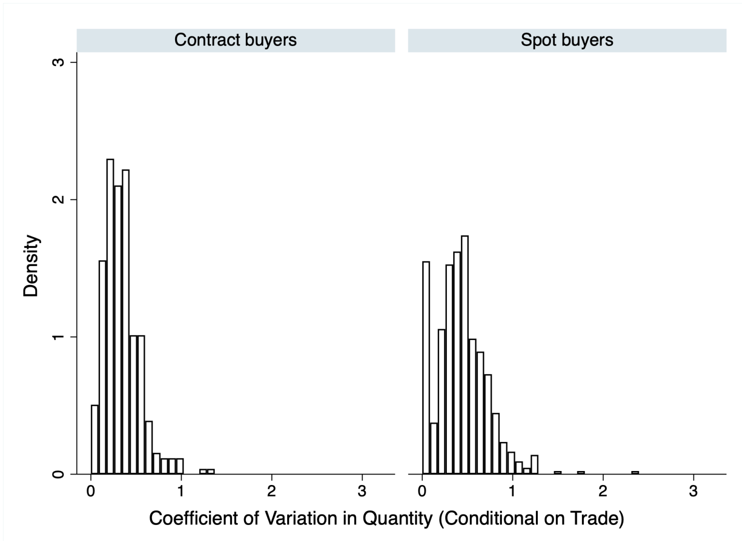

Notes. The coefficient of variation equals the standard deviation of quantity over the mean quantity. Months where no trade occurs are excluded in the calculation of standard deviation and mean. Each observation is a buyer within a region, fiber, and year. There is a large density of spot buyers with a coefficient of variation equal to zero because some buyers trade the exact same quantity every month they trade.

Online Appendix B Details on Estimation

In this section, we provide further details on estimation of the model. Estimation occurs in two steps.

Step 1: Spot Buyer Sample. Step one uses the spot buyer sample to estimate the parameters that govern the rebate and quantity type distribution for spot buyers as well as the extent of trading frictions: $(\beta_\theta, \beta_Y^S, \sigma_{\theta\theta}, \sigma_{YY}^S, \sigma_{Y\theta}, \gamma)$. We use two-stage generalized method of moments. There are six moments:

$$(R_{ijt} - \mathbb{E}[\theta_{ijt}|\tau_{ijt} = 1, X_{ijt}])\, X_{ijt}\tau_{ijt}$$
$$(q_{ijt} - \mathbb{E}[Y_{ijt}|\tau_{ijt} = 1, X_{ijt}])\, X_{ijt}\tau_{ijt}$$
$$\big((R_{ijt} - \mathbb{E}[\theta_{ijt}|\tau_{ijt} = 1, X_{ijt}]])^2 - \mathrm{Var}(\theta_{ijt}|\tau_{ijt} = 1, X_{ijt}])\big)\, \tau_{ijt}$$
$$\big((q_{ijt} - \mathbb{E}[Y_{ijt}|\tau_{ijt} = 1, X_{ijt}]])^2 - \mathrm{Var}(Y_{ijt}|\tau_{ijt} = 1, X_{ijt}])\big)\, \tau_{ijt}$$
$$(R_{ijt}q_{ijt} - \mathbb{E}[\theta_{ijt}Y_{ijt}|\tau_{ijt} = 1, X_{ijt}])\, \tau_{ijt}$$
$$(\tau_{ijt} - \mathbb{E}[\tau_{ijt}|X_{ijt}])X_{ijt}$$

In order to estimate the model using generalized method of moments, we require the model-predicted first and second moments of rebate and quantity, as well as the model-predicted first moment of the trade indicator. Conditional on trade occurring, the rebate R_{ijt} equals the rebate type θ_{ijt}. Conditional on trade occurring, quantity q_{ijt} equals the quantity type Y_{ijt}. Therefore, once we have an expression for the probability of trade conditional on X_{ijt}, we can calculate the relevant moments.

We use an approximation of the conditional probability of trade $E[\tau_{ijt}|X_{ijt}]$. Monte carlo simulation is computationally infeasible. For notational convenience, we suppress dependence on market j and month t. Trade is determined as follows:

1. The types (θ_i, Y_i) realize.

2. A subset N of S of size $\gamma|S|$ is selected at random.

3. The seller lines up the buyers i in N according to mill gate price $p_i := \overline{p}(1 - \theta_i) - c_i$. Note that p_i is i.i.d. normally distributed across buyers (because we condition on logistics costs c_i in X_i and θ_i is i.i.d. normal).

4. The seller calculates cumulative quantity type $\tilde{Y}_i$ for each buyer according to this ordering.

5. The seller trades a quantity $q_i = Y_i$ for those buyers where cumulative quantity type $\tilde{Y}_i$ is less than or equal to total spot quantity Q^S. All other buyers do not trade.

First, we condition on buyer i having type (θ_i, Y_i) and on buyer i being in the set N. Trade occurs if and only if cumulative quantity type $\tilde{Y}_i$ is at most Q^S:

$$\begin{aligned}
&\mathbb{E}[\tau_i|\theta_i, Y_i, i \in N, X_i] \\
&= Pr(\tilde{Y}_i \leq Q^S|\theta_i, Y_i, i \in N, X_i) && (1) \\
&= Pr\left(\sum_{i' \in N: p_{i'} > p_i} Y_{i'} + Y_i \leq Q^S \middle| \theta_i, Y_i, i \in N, X_i\right) && (2) \\
&= Pr\left(\frac{1}{|i' \in N : p_{i'} > p_i|} \sum_{i' \in N: p_{i'} > p_i} Y_{i'} \leq \frac{1}{|i' \in N : p_{i'} > p_i|}(Q^S - Y_i) \middle| \theta_i, Y_i, i \in N, X_i\right). && (3)
\end{aligned}$$

Equation (1) follows by step five of the algorithm that determines trade. Equation (2) follows by the definition of $\tilde{Y}_i$. Equation (3) follows by rearrangement.

Next, we calculate the conditional distribution of

$$\frac{1}{|i' \in N : p_{i'} > p_i|} \sum_{i' \in N: p_{i'} > p_i} Y_{i'}.$$

Consider the distribution of each $Y_{i'}$ conditional on $p_{i'} > p_i$. Because $(Y_{i'}, p_{i'})$ are jointly normal, it is straightforward to calculate the mean and variance of $Y_{i'}$ conditional on $p_{i'} > p_i$. By the central limit theorem, the average of random variables with known mean and variance approaches a normal distribution with known mean and variance. Let F^k denote this distribution if there are k buyers i' in N with $p_{i'} > p_i$.

Next, we integrate over all values of k. Let p^k denote the random variable equal to the k'th smallest value of p_i among $|N| = \gamma|S|$ draws from the common type distribution. By normality, this random variable (an order statistic) has a known distribution. The distribution is binomial with known mean and variance. We approximate this binomial distribution with a normal distribution G^k that has pdf g^k. Using this knowledge, we conclude:

$$\mathbb{E}[\tau_i|\theta_i, Y_i, i \in N, X_i] \approx \int_{k=1}^{\gamma|S|} F^{k-1}\left(\frac{1}{k-1}(Q^S - Y_i)\middle| Y_i\right) g^k(p^k = p_i|p_i)dk.$$

To complete the derivation, we integrate over the realization of buyer i's types and over the placement of i in N. Let H denote the normal type distribution. Each buyer is placed in N with probability γ that is independent of the buyer's type. Therefore,

$$\mathbb{E}[\tau_i \mid X_i] \approx \gamma \int_{\theta_i, Y_i} \int_{k=1}^{\gamma|S|} F^{k-1}\left(\frac{1}{k-1}(Q^S - Y_i)\middle| Y_i\right) g^k(p^k = p_i|p_i)dkdH(\theta_i, Y_i|X_i).$$

To calculate this expression, we use quadrature rules. In simulation exercises, we are able to recover true parameter values using GMM. This supports our approximation.

Step 2: Contract Buyer Sample. In the second step, we use the contract buyer data to estimate the remaining parameters: $(\beta_Y^C, \beta_\phi, \sigma_{YY}^C, \delta_{ijy})$. We do so using ordinary least squares and logistic regression. First, to estimate β_Y^C and σ_{YY}^C we regress quantity (conditional on trade) on the secondary covariate vector X_{ijt}^C. Second, to estimate β_ϕ, we run a logistic regression of the trade indicator on the primary covariate vector X_{ijt}. Third, to estimate the bargaining parameter δ_{ijy}, we use ordinary least squares. The model implies:

$$\mathbb{E}[\text{TRANSFER, } i \text{ IN CONTRACT}] = (1-\delta_{ijy})\mathbb{E}[\text{BUYER OUTSIDE OPTION}]$$
$$+\,\delta_{ijy}(\mathbb{E}[\text{SELLER SPOT PROFITS, } i \text{ IN SPOT}] - \mathbb{E}[\text{SELLER SPOT PROFITS, } i \text{ IN CONTRACT}).$$

To calculate the terms of this equation, we use monte carlo simulation and the realization of market conditions. For all twelve months of the year, we simulate trade when buyer i is placed in spot and we use the realized outcomes when buyer i is placed in contract. We then estimate the expectation with the average over the twelve months. In the model, all remaining variation is due to the difference between ex-ante expectation and ex-post realizations. By definition, these differences are mean-zero, allowing us to estimate δ_{ijy} with ordinary least squares.