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

Olivier Darmouni
Columbia Business School
omd2109@columbia.edu

Simon Essig Aberg
Harvard Business School
sessigaberg@hbs.edu

Juha Tolvanen
Tor Vergata
juha.tolvanen@uniroma2.it

December 8, 2023

Abstract

Firms in decentralized markets often trade using quantity contracts, agreements that specify quantity prior to the point of sale. These contracts are valuable because they provide quantity assurance, as trading frictions could prevent a buyer and seller from matching in the spot market. However, quantity contracts prevent sellers from optimally allocating production across buyers after market conditions realize. Using proprietary invoice data, we estimate a model of quantity contracts in the pulp and paper industry. The average value of a quantity contract is 10% of profits, but would be 55% larger if sellers could optimally allocate production after market conditions realize.

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, Alessandro Gavazza, Jun Ishii, Cristopher Kingston, Amanda Kowalski, Rocco Macchiavello, Alexander MacKay, Jessica Reyes, Tobias Salz, Charles Smith, Nahim Bin Zahur, as well as seminar participants at the University of Vienna, University of Salzburg, University of Amsterdam, University of Manchester, Università Cattolica del Sacro Cuore, Tor Vergata, and conference participants at EARIE 2021, EEA-ESEM 2021 and EWMES 2021 for their helpful feedback. Essig Aberg: this material is based upon work supported by the National Science Foundation Graduate Research Fellowship Program.

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 as a way to increase trade.[1] However, intermediation is not without costs and may even reduce welfare.[2] In general, alternative structures can arise to mitigate the costs of trading frictions.

In this paper we show that quantity contracts, agreements which specify quantity in advance of trade, are valuable in decentralized markets because they reduce the costs of trading frictions. Quantity contracts may allow prices to be set even after delivery. Many industries use quantity contracts, including coal [Joskow, 1988], coffee [Blouin and Macchiavello, 2019], dry bulk shipping [Brancaccio et al., 2020], and liquefied natural gas [Zahur, 2022].

We argue the value of a quantity contract is a combination of an assurance value and a lost option value. Quantity contracts assure that a certain quantity will be traded. Quantity assurance is valuable in decentralized markets because trading frictions could prevent a buyer and seller from matching in the spot market. There are many reasons buyers and sellers may not match in the spot market, and we refer to these collectively as 'trading frictions.' For example, it is difficult to organize the logistics of transporting a bulky commodity between global trading partners. There may also be transaction costs associated with writing a spot contract. Quantity contracts reduce such transaction costs associated with organizing trade in decentralized markets. 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 is stable over time.

The downside of the use of quantity contracts as quantity assurance is a lost option value. Sellers cannot freely allocate production across buyers after non-contractible market conditions realize. When a seller faces an unexpectedly strong spot market, quantity contracts restrict its ability to reallocate quantity from contract buyers to spot buyers. Conversely, contracts prevent reallocation from spot buyers to contract buyers when the spot market is weak.

Our main contribution is to develop an empirical model of quantity contracts and spot trade and use it to quantify the value of quantity contracts in the pulp and paper industry. The pulp industry is well-suited for analysis of quantity contracts. Pulp is a homogeneous good up to well-defined grades, yet 83% of trade

[1] See, for example, Gavazza [2016], Farboodi et al. [2019], Cuesta et al. [2019], Robles-Garcia [2020], and Gavazza and Lizzeri [2021].

[2] As in Hsieh and Moretti [2003], Leslie and Sorensen [2014], and Gavazza [2016].

occurs through annual quantity contracts instead of on the spot. The homogeneity of pulp suggests that quantity contracts do not primarily serve to prevent moral hazard on product quality (as in Lambert [1983]). Furthermore, prices vary significantly both across months and within a month, even holding fixed buyer and seller characteristics. Quantity contracts typically use indexed prices instead of fixed prices, suggesting they do not primarily serve to hedge against price risk (as in Wolak [2000]).

We use proprietary invoice data from a large pulp producer from 2014 to 2019. In addition to the price and quantity of each contract and spot transaction, we observe contract fulfillment rebates, logistics costs, and production costs. We combine these data with data on regional gross price indices and average rebates. The detailed breakdown of costs, prices and rebates gives us a clear picture of profit margins, pairwise quantities, and frequency of trade, while the global variables allow us to control for market-level shocks. The seller we observe is among the largest across products and regions. Its representatives confirmed that its trading behavior is typical of other large firms in the industry. Prior work [Marshall, 2020] has used comparable granular data from a single firm to study search in decentralized markets.

We build a simple model of decentralized markets and quantity contracts. The model explains which buyers sign contracts based on the value of the contract relative to spot trade and a self-enforcement constraint.[3] Contracted prices and quantities depend on a set of contractible variables. Prices and quantities are negotiated via 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. This model of negotiated prices is similar to the approach in Allen et al. [2019].

Each month, a seller allocates inelastic production across contract and spot buyers.[4] Heterogeneous buyers have downward sloping residual demand curves. The seller first trades with the contract buyers according to the quantity contracts. The seller then matches with a subset of the spot buyers and trades its remaining quantity with them. The seller makes take-it-or-leave-it price offers to spot buyers. Matches occur randomly and the match probability depends on a set of observable buyer characteristics. The set of spot buyers that the seller matches with is non-contractible.

The model admits a decomposition of the value of quantity contracts into a weighted average of the assurance value and the lost option value. The assurance value is the difference in payoffs when the buyer trades the contracted quantity with the seller versus when the buyer does not trade with the seller at all. The lost option value is the difference in payoffs when the buyer trades the ex-post optimal quantity versus

[3] We do not microfound the self-enforcement constraint. A large literature has studied the consequences of limited enforcement in similar settings [Macchiavello, 2022]. In our setting, the use of indexed pricing mitigates the threat of opportunism [Blouin and Macchiavello, 2019], as does the homogeneity of pulp. Moreover, contracts typically include a 'contract fulfillment rebate' that is paid if the terms of the contract are satisfied at the end of the year. In our data, we do not observe any instances of default. In Online Appendix C.2, we present an empirical exercise that does not find evidence of hold-up in our setting.

[4] Empirically, the elasticity of production with respect to market price is indistinguishable from zero. Large fixed costs of production lead sellers to maximize capacity utilization in this industry.

when the buyer trades the contracted quantity. Both are non-negative in expectation. The weight on the assurance value is the probability that the buyer would not match with the seller on the spot. Therefore, contracts only create value when trading frictions are sufficiently large.

We evaluate predictions of the model in the invoice data. Consistent with trading frictions, spot buyers trade far less frequently than contract buyers. The average spot buyers trades in 55% of months and the average contract buyer trades in 90% of months. This difference holds up when only considering buyers that switch between spot and contract status: the frequency of trade is 17 percentage points higher when a buyer trades with a contract versus when that same buyer trades on the spot.

The model predicts that more quantity is contracted for buyers with lower logistics costs and for whom the spot match probability is lower. In line with these predictions, spot buyer logistics costs are 16% higher than contract buyer logistics costs. Additionally, in markets where spot buyers trade more frequently, less quantity is contracted.[5] The model admits a test for the presence of a lost option value: on average, the minimum spot margin should equal the minimum contract margin, but these will not be equal with probability one if there is a lost option value. Empirically, the average minimum contract price is within one percent of the average minimum spot price, but the standard deviation of the price difference is 12% of the average.

To quantify the value of quantity contracts, we propose an empirical specification of the model. In the specification, residual demand curves are piece-wise constant. We assume this functional form because empirically there is no detectable correlation between logistics costs and quantity even though there is a strong correlation between logistics costs and price conditional on buyer fixed effects. We allow for heterogeneity in residual demand curves and in spot match probabilities by, among other variables, buyer size and an indicator of whether the buyer ever signs a contract during the sample period. It is important to allow for heterogeneity in past contract status and size because contract buyers and large buyers may have different spot match probabilities than the typical spot buyer. We also allow for heterogeneity in the selection of contract buyers by the same set of buyer characteristics.

We prove the joint distribution of prices, quantities, trade patterns, and covariates identify the model parameters. The correlation between spot buyer trade probabilities and profit margins identifies the spot match probability. Higher spot match probabilities lead the seller to be more selective when choosing trading partners, increasing this correlation. The distribution of spot buyer prices when the seller's spot market is unexpectedly weak identifies the distribution of residual demand curves: the seller makes take-it-or-leave-it price offers to spot buyers, and it must trade with all spot buyers it matches with when the spot market is

[5]There may be market-level implications of this dynamic. Harris and Nguyen [2023] study how contracts increase search costs at the market level because they reduce spot market thickness.

unexpectedly weak. Contract buyer prices relative to contract values identifies buyer 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 spot match probabilities and residual demand curves. The second step uses regressions on the sample of contract buyers to estimate other parameters of the residual demand curves, the contract buyers' bargaining strength relative to the seller, and the parameters that govern the selection of contract buyers. We estimate a contract buyer relative bargaining strength of 29%, consistent with the fact that sellers are much larger than buyers in the industry. Furthermore, the estimated average spot match probability is 66%, which is larger than the observed spot trade probability of 55%. Therefore, the seller is somewhat selective about the spot buyers it trades with, generating a positive correlation between spot buyer trade probabilities and profit margins. However, the seller is not perfectly selective. A 95% confidence interval around our estimated matching probability excludes 100%, evidence that trading frictions are sizeable.

Using the estimated parameters, we quantify the value of quantity contracts. The average value of a quantity contract is 10% of profits, defined as margin times quantity. The value is positive for 82% of contracts. Sampling variation could explain some negative values. Quantity contracts increase profits by 29% relative to a counterfactual where all buyers trade on the spot. We decompose this value into an assurance value and lost option value. We find that the lost option value is sizable—the mean value of a contract would be 15% of profits if the seller could optimally reallocate quantity after market conditions realize.

We quantify the effects of trading frictions and logistics costs on contract values. A 10% increase in the probability of matching with the seller on the spot causes a 26% reduction in the average value of a quantity contract. Because the value of contracts decreases, the fraction of quantity that is contracted falls from 68% to 54%. Next, a 10% increase in logistics costs causes a 3% reduction in the median value of a quantity contract. This result is consistent with the model prediction that the value of a quantity contract increases in a buyer's profitability.

We evaluate the robustness of the model to two assumptions. First, we consider an alternative pricing model where the seller makes take-it-or-leave-it price offers to contract buyers, instead of bargaining over price via a Nash-in-Nash protocol. This alternative pricing model allows us to identify residual demand curves separately for spot and contract buyers. Second, we allow the seller to use inventory to smooth sales. In both cases, our results are largely unchanged.

We then consider two extensions to the model. First, we allow quantity contracts to reduce the variance of quantity conditional on trade. Seller profits are concave with respect to quantity traded with a contract buyer. The first unit sold on the spot receives a higher price than later units, so the marginal profit of spot trade decreases in total quantity traded on the spot. Consequently, sellers value reliable, low variance, contract trade.[6] Quantity contracts may increase buyer reliability if buyers can smooth their demand with storage. Empirically, the coefficient of variation of spot buyer quantity is 9% higher than the coefficient of variation of contract buyer quantity.[7] If contracts reduce the coefficient of variation in contract buyer quantity by 9%, then the average value of a quantity contract increases to 27% of profits.

Second, we investigate how quantity contracts affect price dispersion. Prior work finds that trading frictions [Chandra and Tappata, 2011, Kaplan and Menzio, 2015, Stango and Zinman, 2016] and price discrimination [Zahur, 2023] generate price dispersion in decentralized markets. On one hand, quantity contracts may reduce price dispersion because sellers insure against the necessity of trade with low profit margin spot buyers. On the other hand, quantity contracts may increase price dispersion because sellers are more exposed to non-contractible inter-month variation in spot market strength. We find that these two effects wash out and quantity contracts reduce price dispersion by 3% on average, though in some markets contracts reduce price dispersion by as much as 27%.

Throughout, we limit the scope of our analysis by focusing on bilateral contracting instead of market equilibrium because we use data from one firm. Our notion of surplus holds fixed the contracting behavior of other buyers and sellers in the market, which may have implications for equilibrium contract quantities. Our relationship-level analysis of contracts complements the market-level analysis in Harris and Nguyen [2023].

The results of our analysis have general relevance. An emerging empirical literature studies industries where firms agree to quantities prior to the point of sale even when spot trades are available. Examples include liquefied natural gas, coffee, roses, truckload freight, coal, beef processing, and garments. We complement existing explanations with the finding that quantity contracts can mitigate the cost of trading frictions. This explanation is likely to be broadly relevant because a large literature documents the salience of trading frictions in explaining global trade patterns. Furthermore, our model provides a quantifiable micro-foundation for the popular assumption in the literature on quantity and forward contracts that firms value reliability.

1.1 Related Literature

First, we contribute to the literature on inter-firm contracting and relationships.[8] Zahur [2022] finds that long-term contracts prevent hold-up on capital investment and estimates a structural model of contracts and spot trade to quantify these effects in the market for liquefied natural gas. We contribute by providing a

[6]The finding that sellers endogenously value reliability complements existing work where the value of reliability arises exogenously from a kink in payoff functions [Macchiavello and Morjaria, 2015] or where reliability may enter in contract surplus as a reduced-form parameter [Zahur, 2022, Harris and Nguyen, 2023].

[7]The coefficient of variation equals the standard deviation divided by the mean.

[8]See Lafontaine and Slade [2013] for a review.

complementary mechanism through which contracts create value in a similar setting. We also incorporate trading frictions into our model of spot trade, which are crucial to understanding the value of contracts in our setting. Our analysis also complements Harris and Nguyen [2023], who find that contracts in the truckload freight industry impose a negative externality by reducing spot market thickness and increasing spot search costs. Because we observe the full distribution of spot and contract prices, we are able to identify how contracts reduce such search costs for individual market participants. Consequently, our analysis suggests a feedback effect that may amplify this externality: when spot markets are thin, the value of contracts is higher, which could further reduce spot market thickness.

Another strand of this literature empirically studies relational contracts that rely on dynamic enforcement [Harris and Nguyen, 2021, Cajal-Grossi et al., 2023].[9] We build most directly on the work of Macchiavello and Morjaria [2015] and Macchiavello and Miquel-Florensa [2018], who test empirical predictions of theories regarding relational contracting in the rose and coffee industries. We estimate a structural model of similar contracting relationships to quantify and decompose the value of contracts. The model allows for dynamic enforcement considerations without providing a microfoundation.

Second, we contribute to the literature on search costs and trading frictions in decentralized markets.[10] Allen et al. [2019] and Marshall [2020] study search and pricing in decentralized markets. Gavazza [2016] and Salz [2022] quantify how intermediation affects welfare in markets with search frictions. We introduce quantity contracts to a model of search in a decentralized market. We use our rich data to incorporate heterogeneity into the distributions of both trading friction exposure and buyer outside options.

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.[11] 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.

[9] See Baker et al. [2002] and Li and Matouschek [2013] for theoretical analyses.
[10] See Gavazza and Lizzeri [2021] for a review.
[11] The statistics in this section are from internal presentations by a large seller in the industry.

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.[12]

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 long run, production responds to market conditions through capacity expansions and the construction and shut-down of mills. Nevertheless, there is still monthly variation in production due to planned maintenance, unplanned breakages, and worker strikes. In response to the inelasticity of production in the short run, pulp producers use inventory. However, inventory is constrained by storage limitations (pulp is relatively bulky) 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. As a result, variation in a buyer's suppliers' production is one force that causes variation in its willingness to pay and its residual demand over time.

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. Contracts often include clauses that preclude the resale of pulp, preventing buyers with different valuations from exploiting arbitrage opportunities.

Gross prices are either indexed or negotiated each month. In most contracts, the gross price is indexed.[13]

[12]Specifically, we consider bleached chemical pulp, which is by far the largest segment of chemical pulp.

[13]They are hence similar to the index-priced forward contracts for coffee described in Blouin and Macchiavello [2019]. In

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. The indices combine data on spot and contract transactions. 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. In Figure OA.1, we plot the timeseries of gross price for these contracts and find little difference by price-setting method. If anything, the indexed price timeseries is lagged relative to the negotiated price timeseries. In line with this finding, 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. Because we focus on the static value of quantity contracts, we abstract from the distinction between indexed and negotiated price in this paper.

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.

As with any insurance contract, limited enforcement of quantity contracts is a concern. The main enforcement concern is that the market will change unexpectedly and the buyer will refuse to trade with the seller or vice versa. The use of indexed pricing instead of fixed pricing mitigates but does not eliminate counterparty risk [Blouin and Macchiavello, 2019]. 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 are typically less than 2%. 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.[14] Dynamic considerations may also prevent short-term opportunism [Harris and Nguyen, 2021]. Many buyers and sellers have relationships dating back to the 1980s or earlier, so preserving reputation may be important. We conduct an empirical analysis in Online Appendix C.2 and are unable to find evidence of hold-up.

Other industries feature comparable quantity contracting. In some of them, contract enforcement is a greater concern. For instance, Macchiavello and Morjaria [2015] study limited enforcement in a setting with variable product quality, Zahur [2022] studies hold-up in a setting with capital-intensive investment decisions, and Harris and Nguyen [2021] analyze how dynamic incentives prevent short-term opportunism when contracts fix price. Because enforcement is less of a concern in our setting, we are able to focus on the

contrast, none of the quantity contracts in our data fix prices at the time of contracting as in the truckload freight industry [Harris and Nguyen, 2021, 2023].

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

role of quantity contracts in mitigating the costs of trading frictions.

Structure 2: Spot Trade. Trade that does not occur via quantity contracts occurs at a decentralized spot market. 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. Three trading frictions may inhibit efficient spot trade. First, search frictions in the global shipping industry may prevent trade from occurring. Even if a seller finds a high surplus trading partner on the spot, it may not be able to find a ship that can transport pulp due to the spatial misallocation of ships. Brancaccio et al. [2023] find that the social costs of such spatial misallocation are large in the dry bulk shipping industry. Though we do not have data on transportation in our setting and do not model spatial misallocation, our model and findings are consistent with this friction. Second, the decentralized nature of the spot trade introduces search frictions. Spot deals are typically negotiated by buyer and seller side sales managers in real time based on buyers' stochastically varying needs and sellers' availability. Due to the decentralized, real-time nature of the trade, not all of the offers are necessarily on the table when decisions on either side of the market are made. Similarly, as calling and negotiating with buyers takes time, a competitor may snatch an available deal before our firm learns about it, let alone has enough time to compare it to its other options. Furthermore, different sales managers within a firm do not necessarily have a complete picture of all of the available deals when deciding on which of their private offers to accept. These types of search and information frictions can lead to inefficient trades relative to the first-best matches between buyers and sellers. Third, 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. Allen [2014] studies information frictions of this sort and finds them to be quantitatively important in explaining price dispersion.

Since 2020 (after our analysis period), a large physical futures exchange has developed in Shanghai, though its success is far from certain as 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.

Spot trade is a common organizational structure across industries. The spot markets for liquefied natural gas [Zahur, 2022] and truckload freight [Harris and Nguyen, 2023] are quite similar. Some industries have slightly different structures. For instance, the spot market for roses operates through a centralized exchange [Macchiavello and Morjaria, 2015].

Structure 3: Vertical Integration. Of the ten largest pulp buyers, seven are partially vertically integrated.[15] 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 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 market price and average rebate.[16] We restrict analysis to transactions between non-integrated mills.

We merge the seller's invoices with its 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 only 92 write contracts for at least one product. Despite this imbalance, 83% of quantity is contracted.

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

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

[16] For market price, we use an internal estimate of the market-wide quantity-weighted average price across contract and spot buyers. Publicly available indices of market price are released with a month lag.

Table 1: Descriptive Statistics

	N	Mean	SD/Mean	Median	IQR
Panel A: Buyer characteristics.					
Contract buyer	268	0.34			
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)	137	24,030	1.89	9,167	16,294
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 (€/ton)	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		32,134
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.

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.

The categorization of buyers varies over fibers, regions, and years. Annual contracts typically include separate price and quantity clauses across regions and fibers. 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 (which are paid by sellers in this industry), we arrive at the mill gate price. We refer to mill gate price as the ‘margin’. One representative of the seller noted that the seller aims to maximize trade at a high margin without considering production costs that are linked to invoices. This is true because the seller considers production costs as largely fixed due to the short-run inelasticity of production. Therefore, we use the margin to measure the seller’s flow profits.

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 Model

This section presents a stylized model of a decentralized market. The primary purpose of the model is to demonstrate how quantity contracts can be valuable when spot trade is subject to search frictions. The model also generates a decomposition of the value of quantity contracts, comparative statics, and other testable predictions.

We present the model from the perspective of a single large seller because that is the context of our empirical analysis. Consequently, the notion of ‘value’ is bilateral. The purpose of the model and empirical exercise is to explain the use of contracts, not to measure the net market-wide welfare effects. We do not measure the externality that signing a contract imposes on other sellers by reducing the contract buyer’s residual demand.

3.1 Environment

A large seller produces quantity Q_{jt} in market j and month t. We assume production is determined exogenously. It does not depend on short-term market conditions or the set of buyers with contracts. There are substantial fixed costs involved in production. Table OA.1 shows that there is no detectable correlation between production and market price. In Section 7.3, we consider an extension where the seller can use inventory to respond to short-term market conditions and the results are largely unchanged.

Buyer i in market j has residual inverse demand curve $\theta_{ijt}(q)$ in month t. We assume this function is non-negative and weakly decreasing in quantity transacted with the seller q. The residual inverse demand curve equals the marginal value to the buyer of the qth unit bought from the seller. We want to emphasize that variation in the residual inverse demand curve can arise from many sources: downstream demand for the buyer’s paper products, the buyer’s transactions with other sellers, the buyer’s use of inventory, and heterogeneity in buyer bargaining power. We remain agnostic to the specific microfoundation of $\theta_{ijt}(q)$, though in general this microfoundation is important for a market-wide welfare analysis of quantity contracts.

The seller trades a quantity $q_{ijt} \geq 0$ at price $p_{ijt} \geq 0$ with buyer i in market j and month t. The variable $\tau_{ijt} \in \{0,1\}$ indicates whether trade occurs: $\tau_{ijt} = 1[q_{ijt} > 0]$. The marginal cost of trade is constant and equals $c_{ijt} \geq 0$. The seller's profit in market j and month t is

$$\Pi_{jt} = \sum_i (p_{ijt} - c_{ijt}) q_{ijt}.$$

The buyer's payoff is the area under the residual inverse demand curve net of price:

$$U_{ijt} = \int_0^{q_{ijt}} \theta_{ijt}(q) - p_{ijt} dq.$$

Buyers are classified as either spot buyers or contract buyers each year. Let S_{jy} denote the set of spot buyers in market j and year y and let C_{jy} denote the set of contract buyers in market j and year y. The game takes place in three stages. First, at the start of the year, the set of contract buyers are selected. Second, each month, the seller trades with contract buyers. Third, each month, the seller trades with spot buyers. We describe each of these three stages, starting in stage 3 and working backwards.

3.2 Spot Trade

As a consequence of trading frictions, the seller is unable to trade with all spot buyers in S_{jy}. Let $m_{ijt} \in \{0,1\}$ indicate whether buyer $i \in S_{jy}$ matches with the seller in month t. The key parameter that indexes the degree of trading frictions is the spot match probability

$$\gamma_{ijt} = \mathbb{E}_{ijt} m_{ijt},$$

where the subscript ijt on the expectation operator indicates that there may be heterogeneity in the spot match probability over buyers i, markets j, and months t. In the empirical specification, we allow for heterogeneity in γ_{ijt} by observable buyer characteristics including buyer size and whether the buyer signs a contract during the sample period. We assume that matches occur independently across buyers: m_{ijt} is distributed i.i.d. Bernoulli with parameter γ_{ijt}.

Conditional on the set of buyers with m_{ijt} and residual demand curves $\theta_{ijt}(q)$, the seller makes take-it-or-leave-it price and quantity offers to maximize profits. This assumption implies that all buyers will receive zero payoff in equilibrium. Consequently, this assumption can be thought of as a normalization of the residual demand curves $\theta_{ijt}(q)$. Any spot buyer bargaining power can be captured by the residual demand curve $\theta_{ijt}(q)$.

Suppose the seller traded Q^C_{jt} with contract buyers such that $Q^S_{jt} = Q_{jt} - Q^C_{jt}$ of its production remains for the spot buyers. Spot buyer prices and quantities are given by the solution to the following problem:

$$\Pi^S_{jt}(S_{jy}, Q^S_{jt}) = \max_{(p_{ijt}, q_{ijt})_{i \in S_{jy}}} \sum_{i \in S_{jy}} m_{ijt}(p_{ijt} - c_{ijt})q_{ijt} \tag{1}$$

$$\text{subject to } U_{ijt} \geq 0 \text{ for all } i \in S_{jy}$$

$$\text{and } \sum_{i \in S_{jy}} m_{ijt} q_{ijt} \leq Q^S_{jt}.$$

The first set of constraints are the spot buyer participation constraints. The second constraint is the seller's production constraint. The spot profit function $\Pi^S_{jt}(S_{jy}, Q^S_{jt})$ depends explicitly on the set of spot buyers S_{jy} and the total spot quantity Q^S_{jt} as these will be important for the valuation of quantity contracts.

Because the residual demand curves are positive and weakly decreasing, the spot buyer participation constraints will be binding, so in equilibrium

$$m_{ijt} p_{ijt} q_{ijt} = \int_0^{q_{ijt}} \theta_{ijt}(q) dq$$

and U_{ijt} equals zero.

3.3 Contract Trade

Contracts are maps from contractible variables to price and quantity outcomes. Let X^C_{ijt} denote the set of contractible variables. For example, in our empirical analysis we will specify that X^C_{ijt} includes variables such as market price and logistics costs, but does not include the idiosyncratic demand realizations of other buyers or the set of spot buyers which match with the seller. The price and quantity maps are negotiated at the start of the year between a contract buyer and the seller and are denoted $p_{ijy}(X^C_{ijt})$ and $q_{ijy}(X^C_{ijt})$.

The price and quantity maps are determined through a Nash-in-Nash bargaining protocol conditional on X^C_{ijt}. Therefore, quantity is chosen to maximize the expected pairwise surplus and price is chosen to split the surplus based on a bargaining weight. We assume that if contract negotiations fail, then the buyer joins the pool of spot buyers S_{jy}. We discuss the issue of enforceability in the following sub-section.

The value of a quantity contract equals the difference in expected payoffs when the buyer signs a contract versus when the buyer trades on the spot. Formally, the seller's contract profits are given by

$$\Pi^C_{jt} = \sum_{i \in C_{jy}} (p_{ijy}(X^C_{ijt}) - c_{ijt}) q_{ijy}(X^C_{ijt})$$

and the buyer's contract payoff is given by

$$U^C_{ijt} = \int_0^{q_{ijy}(X^C_{ijt})} \theta_{ijt}(q) - p_{ijy}(X^C_{ijt}) dq.$$

The ex-post contract value is

$$W_{ijt} = U^C_{ijt} + \Pi^C_{jt} + \Pi^S_{jt}\left(S_{jy}, Q_{jt} - \sum_{i' \in C_{jy}} q_{i'jy}(X^C_{i'jt})\right) - \Pi^S_{jt}\left(S_{jy} \cup \{i\}, Q_{jt} - \sum_{i' \in C_{jy} \setminus \{i\}} q_{i'jy}(X^C_{i'jt})\right)$$

because the buyer obtains zero payoff when it trades on the spot. The ex-post contract value does not depend on the transfer $p_{ijy}(X^C_{ijt}) q_{ijy}(X^C_{ijt})$. The ex-ante contract value is

$$V_{ijy} = \mathbb{E}_{ijy} W_{ijt}.$$

The primary objective of the empirical analysis is to quantify the contract value V_{ijy}.

The price and quantity maps are selected to solve the Nash product conditional on all the other contracts that have been signed and X^C_{ijt} subject to the constraints that both factors in the product are non-negative.[17] The buyer's bargaining power in the contract negotiations is $\delta_{ijy} \in [0, 1]$.[18] Assuming that there exists a contract that satisfies the participation constraints, the quantity map is the solution to

$$q_{ijy}(X^C_{ijt}) \in \arg\max_q \mathbb{E}_{ijy}[W_{ijt} | X^C_{ijt}]. \tag{2}$$

[17]The Nash product is

$$\mathbb{E}_{ijy}[U^C_{ijt}]^{\delta_{ijy}} \mathbb{E}_{ijy}\left[\Pi^C_{jt} + \Pi^S_{jt}\left(S_{jy}, Q_{jt} - \sum_{i' \in C_{jy}} q_{i'jy}(X^C_{i'jt})\right) - \Pi^S_{jt}\left(S_{jy} \cup \{i\}, Q_{jt} - \sum_{i' \in C_{jy} \setminus \{i\}} q_{i'jy}(X^C_{i'jt})\right)\right]^{1-\delta_{ijy}}.$$

[18]The buyer may received additional rents by bargaining in the spot market, but those are normalized to zero.

The transfer is given by

$$
\begin{aligned}
p_{ijy}(X^C_{ijt})q_{ijy}(X^C_{ijt}) =& \delta_{ijy}\mathbb{E}_{ijy}\Bigg[- c_{ijt}q_{ijy}(X^C_{ijt}) + \Pi^S_{jt}\Bigg(S_{jy}, Q_{jt} - \sum_{i' \in C_{jy}} q_{i'jy}(X^C_{i'jt})\Bigg) \\
&\qquad - \Pi^S_{jt}\Bigg(S_{jy} \cup \{i\}, Q_{jt} - \sum_{i' \in C_{jy} \setminus \{i\}} q_{i'jy}(X^C_{i'jt})\Bigg) \mid X^C_{ijt}\Bigg] \\
&+ (1-\delta_{ijy})\mathbb{E}\left[\int_0^{q_{ijy}(X^C_{ijt})} \theta_{ijt}(q)dq | X^C_{ijt}\right].
\end{aligned} \tag{3}
$$

If the buyer has all the bargaining power ($\delta_{ijy} = 1$), then the transfer will be such that the seller is indifferent between the buyer trading with a contract and trading on the spot. If the seller has all the bargaining power ($\delta_{ijy} = 0$), then the buyer receives zero surplus.[19]

3.4 Contract Buyer Selection

Ex-post, the buyer and seller may have an incentive to deviate from the price and quantity stipulated in the contract. For example, if the set of spot buyers is unusually advantageous to the seller in some month t and this is not contractible, then the seller has an incentive to reduce quantity supplied to contract buyers. Similarly, if the buyer's residual demand curve is unexpectedly low in some month and this is not contractible, then the buyer has an incentive to reduce quantity demanded through the contract. Because we do not observe any defaults empirically, we assume that contract buyers are selected such that the probability of default is zero. The continuation flow value of a contract is $\delta_{ijy}V_{ijy}$ to buyer i and $(1-\delta_{ijy})V_{ijy}$ to the seller, so as long as V_{ijy} is sufficiently large, the contract will be sustainable. Let V^D_{ijt} denote the maximum of the deviation value to the seller divided by $(1-\delta_{ijy})$ and the deviation value to the buyer divided by δ_{ijy}. Abstracting from the repeated nature of the interaction for simplicity, the contract is enforceable if

$$V_{ijy} \geq V^D_{ijt} \text{ with probability 1.}$$

We assume that $V^D_{ijt} \geq 0$ with positive probability, such that this condition implies that $V_{ijy} \geq 0$ for all contract buyers. In equilibrium, the set of contract buyers is the set of buyers where this condition holds:

$$C_{jy} = \{i : V_{ijy} \geq V^D_{ijt} \text{ with probability 1}\}. \tag{4}$$

[19] In Section 7.4, we consider an alternative model where the seller makes take-it-or-leave-it offers, allowing us to incorporate additional heterogeneity into buyer residual demand curves. Our main results are largely unchanged.

This set may be empty and it may not be unique because of inter-dependencies across i in the value of a contract. The set of spot buyers S_{jy} is simply the set of all buyers except those in C_{jy}.

To summarize, an equilibrium in the model is a set of contract buyers C_{jy}, a set of price and quantity maps $(p_{ijy}(X^C_{ijt}), q_{ijy}(X^C_{ijt}))_{i\in C_{jy}}$, and spot buyer prices and quantities $(p_{ijt}, q_{ijt})_{i\in S_{jy}}$ such that:

1. The set of contract buyers C_{jy} is characterized by (4).

2. For contract buyers $i \in C_{jy}$, price p_{ijt} equals $p_{ijy}(X^C_{ijt})$ and quantity q_{ijt} equals $q_{ijy}(X^C_{ijt})$. Contract quantity maps $q_{ijy}(X^C_{ijt})$ are determined by (2) and contract price maps $p_{ijy}(X^C_{ijt})$ are determined by (3).

3. For spot buyers $i \in S_{jy}$, price p_{ijt} and quantity q_{ijt} are a solution to (1).

3.5 The Value of Quantity Contracts

In this section, we decompose the value of quantity contracts into an assurance value and a lost option value. We then show that the value of a quantity contract decreases in logistics costs c_{ijt} and in spot match probability γ_{ijt}. We provide a test that the lost option value equals zero.

The assurance value is the difference in expected payoffs when the buyer trades the contracted quantity versus when the buyer does not trade with the seller:

$$\begin{aligned} V^A_{ijt} =& \mathbb{E}\left[\int_0^{q_{ijy}(X^C_{ijt})} \theta_{ijt}(q) - c_{ijt}dq + \Pi^S_{jt}\left(S_{jy}, Q_{jt} - \sum_{i'\in C_{jy}} q_{i'jy}(X^C_{i'jt})\right)\right] \\ &- \mathbb{E}\left[\Pi^S_{jt}\left(S_{jy}\cup\{i\}, Q_{jt} - \sum_{i'\in C_{jy}\setminus\{i\}} q_{i'jy}(X^C_{i'jt})\right)\middle| m_{ijt} = 0\right]. \end{aligned}$$

The first term equals the sum of buyer and seller payoffs conditional on the buyer trading $q_{ijy}(X^C_{ijt})$. When the buyer trades nothing, it is equivalent to a configuration where buyer i is added to the set of spot buyers and buyer i does not match with the seller on the spot market. Hence the second term equals the sum of buyer and seller payoffs conditional on the buyer trading nothing.

The lost option value is the difference in expected payoffs when the buyer trades the ex-post optimal quantity and when the buyer trades the contracted quantity . This is given by

$$\begin{aligned} V^O_{ijt} =& \mathbb{E}\left[\Pi^S_{jt}\left(S_{jy}\cup\{i\}, Q_{jt} - \sum_{i'\in C_{jy}\setminus\{i\}} q_{i'jy}(X^C_{i'jt})\right)\middle| m_{ijt} = 1\right] \\ &- \mathbb{E}\left[\int_0^{q_{ijy}(X^C_{ijt})} \theta_{ijt}(q) - c_{ijt}dq + \Pi^S_{jt}\left(S_{jy}, Q_{jt} - \sum_{i'\in C_{jy}} q_{i'jy}(X^C_{i'jt})\right)\right]. \end{aligned}$$

The ex-post optimal quantity is the solution to (1) conditional on matching with the seller on the spot market. The seller's profit equals the sum of buyer and seller payoffs because the buyer obtains no payoff on the spot market in equilibrium. Hence the first term equals the sum of buyer and seller payoffs when the buyer trades the ex-post optimal quantity. The second term equals the sum of buyer and seller payoffs conditional on the buyer trading $q_{ijy}(X^C_{ijt})$.

The value of a quantity contract is a weighted average of the insurance value and the lost option value:

$$V_{ijy} = \mathbb{E}_{ijy}[(1-\gamma_{ijt})V^A_{ijt} - \gamma_{ijt}V^O_{ijt}].$$

The weights depend on the spot match probability γ_{ijt}. Both the assurance value and the lost option value are non-negative in expectation. The lost option value is non-negative with probability one because choosing quantity equal to $q_{ijy}(X^C_{ijt})$ for buyer i is in the seller's choice set when it solves (1) conditional on matching with buyer i. The assurance value is non-negative in expectation because the contract selection rule requires $V_{ijy} > 0$ and the weakly positive lost option value enters negatively into the definition of V_{ijy}.

Next, we turn to comparative statics with respect to the spot match probability γ_{ijt} and logistics costs c_{ijt}. Consider a small rightward shift in the distribution of γ_{ijt} of magnitude $d\gamma$. By the envelope theorem, the change in contract value is

$$dV_{ijy} = -\mathbb{E}_{ijy}[V^I_{ijt} + V^O_{ijt}]d\gamma < 0. \tag{5}$$

The value of quantity contracts is larger the spot match probability γ_{ijt} is smaller. To bring this comparative static to the data, we consider a special case where the spot match probability γ_{ijt} is positively correlated withing a market-year. The probability of trade in the spot market is increasing in γ_{ijt} and the value of quantity contracts is decreasing in γ_{ijt}. The contract selection model therefore predicts that fewer buyers sign quantity contracts in markets where the probability of spot trade is larger, all else equal.

Consider a small rightward shift in the distribution of c_{ijt} of magnitude dc. By the envelope theorem, the change in contract value is

$$dV_{ijy} = -\mathbb{E}_{ijy}[q_{ijy}(X^C_{ijt}) - \gamma_{ijt}q^S_{ijt}]dc,$$

where q^S_{ijt} denotes the buyer's counterfactual quantity conditional on matching with the seller on the spot market. This will be negative as long as the spot match probability γ_{ijt} is sufficiently small. The contract selection model therefore predicts that contract buyers have lower logistics costs than spot buyers, all else equal.

Next, we provide a test that the lost option value equals zero. If the lost option value equals zero, then the minimum spot margin equals the minimum contract margin within a market and month. To see this, suppose that the lost option value equals zero and the seller can freely re-allocate quantity across buyers after the market state realizes. Suppose further the minimum spot margin is strictly smaller than the minimum contract margin. Let i^S denote the marginal spot buyer and i^C denote the marginal contract buyer. Consider a perturbation where the seller reduces $q_{i^S jt}$ by dq ind increases $q_{i^C jt}$ by dq. The seller keeps price constant for i^S and reduces the price of i^C by $\delta_{i^C jy}((p_{i^C jy} - c_{i^C jy}) - (p_{i^S jy} - c_{i^S jy}))dq/q_{i^C jt}$. Because buyer i^S has a downward-sloping residual demand curve, buyer i^S is better off and will accept the offer. The seller and buyer i^C split the additional gains created according to (3). The reverse case is analogous. The same argument holds conditional on X^C_{ijt} in an ex-ante sense—the expected minimum spot margin equals the expected minimum contract margin.

We summarize these results in the following proposition:

Proposition 1. *Quantity contracts satisfy the following properties:*

1. *The value of a quantity contract V_{ijy} can be decomposed into an assurance value V^A_{ijt} and a lost option value V^O_{ijt} with weights that depend on spot match probability γ_{ijt}:*

$$V_{ijy} = \mathbb{E}_{ijy}[(1 - \gamma_{ijt})V^A_{ijt} - \gamma_{ijt}V^O_{ijt}].$$

   Both the assurance value and the lost option value are non-negative in expectation.

2. *All else equal, the value of a quantity contract V_{ijy} decreases in the spot match probability γ_{ijt} and in logistics costs c_{ijt}. Fewer buyers sign quantity contracts in markets where the probability of spot trade is larger, all else equal. Contract buyers have lower logistics costs than spot buyers, all else equal.*

3. *Within a market, the expected minimum spot margin equals the expected minimum contract margin. If the lost option value equals zero, then the minimum spot margin equals the minimum contract margin with probability one.*

4 Stylized Facts

In this section we present stylized facts based on the model setup and predictions.

Fact 1: No law of one price. Despite the homogeneity of pulp within well-defined grades, there is sizable heterogeneity in prices and margins conditional on market and month. Figure 1 plots the joint distribution of log price (measured as price after rebates) and log margins (measured as price after rebates

Figure 1: Price and margin dispersion in the spot market

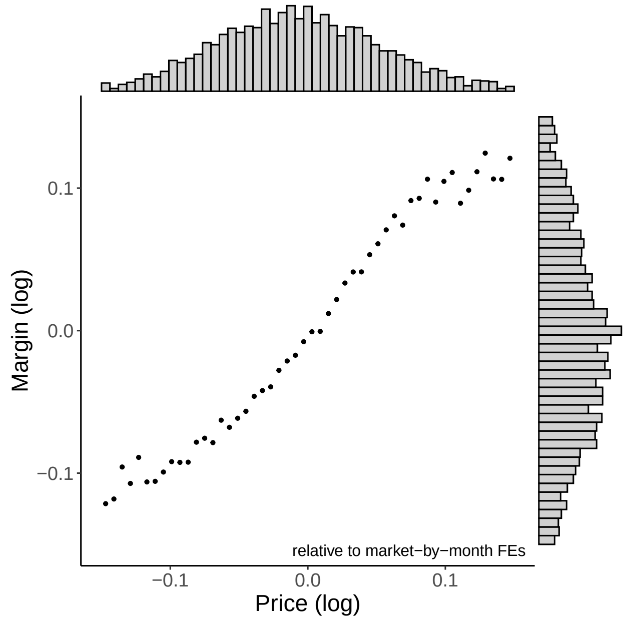

Note. This figure plots the distribution of log price and log margin for spot buyers relative to market-by-month fixed effects. Each points equals the mean of log margin within a bin of log price. The price measure is price after rebates, and the margin measure is mill-gate price, which equals price after rebates minus logistics costs. A market is a combination of fiber and global region. Within a market and year, spot buyers identified as buyers that do not have a performance rebate in their invoices and are not internally prioritized as Tier I by the seller.

minus logistics costs) among spot buyers after removing fixed effects for market-by-month. Relative to the mean within market and month, the standard deviation in log price is 12% and the standard deviation of log margin is also 12%. The existence of price heterogeneity in equilibrium in a homogenous good product market is consistent with a model of trading frictions. Heterogeneity in margins is necessary for contracts to be valuable. Absent margin heterogeneity, the seller would be indifferent between feasible allocations of production across buyers, so contracts would generate neither an assurance value nor a lost option value.

Fact 2: Evidence of trading frictions. Table 2 presents correlations between buyer trade probabilities and contract status. Contract buyers trade in 90% of months and spot buyers trade in 55% of months. The probability of trade is 17 percentage points lower when a buyer trades on the spot versus when that same buyer signs a contract. Most buyers switch from trading on the spot to trading with a contract, as shown in Figure OA.2. There are 12 buyer-markets that switch from trading with a contract to trading on the spot. For all 12, the probability of trade is weakly lower after switching to trading on the spot, as shown in Figure OA.3. For seven of them, the probability of trade is strictly smaller. This result suggests that in counterfactuals where contract buyers trade on the spot, trading frictions would reduce the probability of trade.

Fact 3: Contract buyers have lower logistics costs. The model predicts that the value of a contract

Table 2: Regressions of trade indicator on contract status

	Trade			
Constant	0.899			
	(0.007)			
Spot	-0.353	-0.359	-0.170	-0.181
	(0.009)	(0.008)	(0.045)	(0.046)
Market-month fixed effects		Yes		Yes
Buyer-market fixed effects			Yes	Yes
Observations	9,564	9,564	9,564	9,564
R^2	0.134	0.197	0.386	0.414

Note. Trade occurs at a monthly frequency. A market is a combination of fiber and global region. Within a market and year, spot buyers identified as buyers that do not have a performance rebate in their invoices and are not internally prioritized as Tier I by the seller.

Table 3: Regressions of logistics cost on contract status

	Logistics costs (log)			
Constant	-0.676			
	(0.009)			
Spot	0.164	0.207	-0.081	-0.027
	(0.013)	(0.011)	(0.065)	(0.062)
Market-month fixed effects		Yes		Yes
Buyer-market fixed effects			Yes	Yes
Observations	6,489	6,489	6,489	6,489
R^2	0.023	0.157	0.820	0.838

Note. The sample includes all invoices where trade occurs. A market is a combination of fiber and global region. Within a market and year, spot buyers identified as buyers that do not have a performance rebate in their invoices and are not internally prioritized as Tier I by the seller.

Figure 2: Scatterplot of contract prevalence on spot trade probability

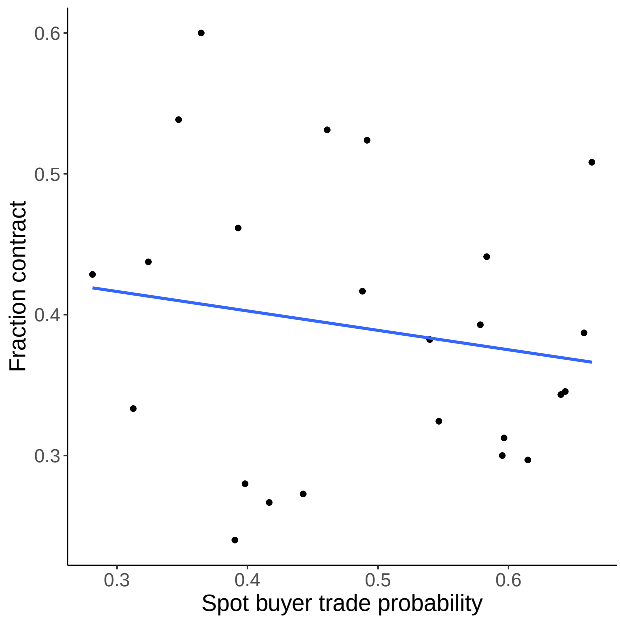

Note. Each observation is a market-year. The spot buyer trade probability is the average of the trade indicator among spot buyers in that market and year. The fraction contracted is the fraction of buyers that sign contracts in that market and year. A market is a combination of fiber and global region. Within a market and year, spot buyers identified as buyers that do not have a performance rebate in their invoices and are not internally prioritized as Tier I by the seller.

decreases in buyer logistics costs, and as a consequence of the contract selection model, higher logistics costs reduce the propensity to sign contracts all else equal. Consistent with this prediction, Table 3 shows that spot buyers have 16% higher logistics costs than contract buyers on average, and this difference rises to 20% upon accounting for market and month fixed effects. There is no detectable difference in logistics costs among buyers that switch between contract and spot status, though there is also less variation in logistics costs within a buyer versus across buyers.

Fact 4: There are more contracts when spot buyer trade probability is low. The model predicts that the value of a contract decreases in the spot buyer match probability. If spot buyer match probability is correlated across buyers within a market and year, then the contract selection model implies that, all else equal, fewer buyers will sign contracts in markets and years where the spot buyer match probability is higher. The spot buyer match probability is unobserved, but is positively correlated with the spot buyer trade probability. Figure 2 plots a (weak) negative correlation between spot buyer trade probability and the fraction of buyers that sign contracts, consistent with this model prediction.

Fact 5: Evidence of the lost option value. A prediction of the model is that, on average, the minimum contract margin equals the minimum spot margin. Figure 3 plots the difference in (log) minimum contract margin and minimum spot margin across markets and month. On average, the two margins differ

Figure 3: Minimum contract margin relative to minimum spot margin

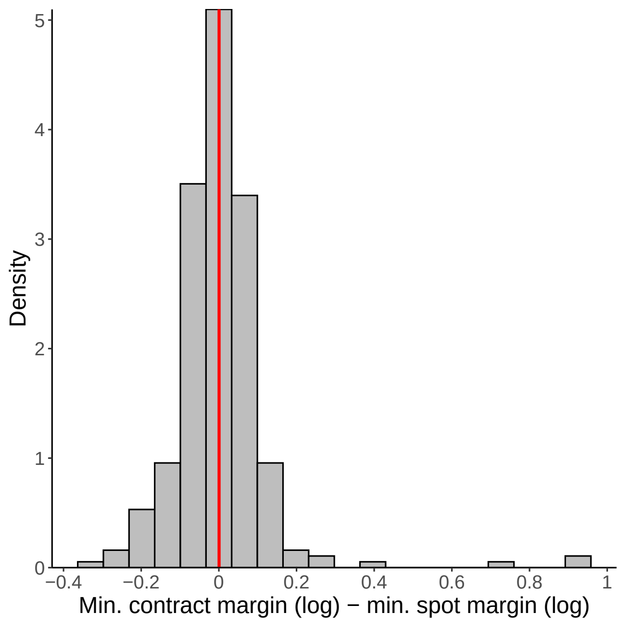

Note. Each observation is the difference between the minimum contract and spot log margins within a market and month. The red line is located at the mean of 0.0005. The margin equals price after rebates minus logistics costs. Within a market and year, spot buyers identified as buyers that do not have a performance rebate in their invoices and are not internally prioritized as Tier I by the seller.

by 0.05%. However, there is significant heterogeneity across markets and month. The average deviation between the two margins is 12%. These deviations are evidence of a lost option value: ex-post, the seller would benefit by reallocating production across buyers.

5 Estimation

In this section, we specify an empirical version of the model in Section 3 and present results on the identification of the model parameters. We then describe the estimation routine and present the estimates.

5.1 Empirical Specification

In line with institutional details that center price negotiations around a rebate off market price, we specify the following functional form for the buyer's residual inverse demand function $\theta_{ijt}(q)$:

$$\theta_{ijt}(q) = \begin{cases} \bar{p}_{jt}(1 - \theta_{ijt}) & \text{if } q \leq Y_{ijt}, \\ 0 & \text{otherwise.} \end{cases}$$

The term $\bar{p}_{jt}$ refers to the market price in market j and month t and θ_{ijt} is the smallest rebate buyer i would be willing to accept for trade of quantity less than Y_{ijt}. We refer to θ_{ijt} as the buyer's 'rebate type' and Y_{ijt} as the buyer's 'quantity type.' A consequence of this functional form is that buyer price-elasticities are trivial. In Table OA.2 , we provide empirical evidence in support of this assumption. There is no detectable correlation between spot buyer quantities and logistics costs conditional on buyer fixed effects, even though there is a large correlation between spot buyer prices and logistics costs conditional on buyer fixed effects.

We parameterize the joint distribution of the rebate type and the quantity type as

$$\begin{aligned}
&\theta_{ijt} = X'_{ijt}\beta_\theta + \epsilon^\theta_{ijt}, \\
&Y_{ijt} = C_{ijy}X^{C\prime}_{ijt}\beta^C_Y + (1 - C_{ijy})X'_{ijt}\beta^S_Y + \epsilon^Y_{ijt}, \text{ where} \\
&\begin{pmatrix} \epsilon^\theta_{ijt} \\ \epsilon^Y_{ijt} \end{pmatrix} | X_{ijt}, X^C_{ijt}, C_{ijy} \overset{iid}{\sim} \mathcal{N}\left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} \sigma_{\theta\theta} & \sigma_{Y\theta} \\ \sigma_{Y\theta} & C_{ijy}\sigma^C_{YY} + (1 - C_{ijy})\sigma^S_{YY} \end{pmatrix} \right).
\end{aligned}$$

In this specification, the covariate vector X_{ijt} includes an indicator for whether the buyer ever signs a contract, a priority indicator, logistics costs, market price, a time trend in months, an indicator that average quantity is greater than 1,000 tons/month, and market fixed effects. The contractible 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_{ijy} indicates whether a buyer i is a contract buyer in market j and month y. We assume that the residuals ϵ^θ_{ijt} and ϵ^Y_{ijt} are distributed i.i.d. normal with conditional mean zero.

For contract buyers, the quantity type Y_{ijt} depends on the contractible covariate vector X^C_{ijt} instead of X_{ijt}. Contract buyers tend to trade much larger quantities than spot buyers, so we include this additional heterogeneity into the quantity type. For the same reason, we allow the variance of ϵ^Y_{ijt} to depend on whether buyer i is a contract buyer. Section 7.4 presents a robustness exercise where we allow the contract buyer rebate type to also depend on a buyer-market-year fixed effect.

An implication of the step functional form is that quantity will equal Y_{ijt} for all buyers that trade, perhaps excluding the marginal buyer in each market. This follows because the marginal value of trade equals the constant $\bar{p}_{jt}(1 - \theta_{ijt}) - c_{ijt}$ for all spot buyers conditional on quantity being at most Y_{ijt}. In the spot market, the solution to (1) can be calculated by ordering buyers in terms of $\bar{p}_{jt}(1 - \theta_{ijt}) - c_{ijt}$ and trading quantity equal to Y_{ijt} until production is exhausted. For contract buyers, there exists an equilibrium where the contracted quantity equals Y_{ijt}, perhaps with the exception of one buyer. This is true because the marginal opportunity cost of trade (lost expected spot revenue) depends on the sum of contracted quantity. Therefore, if the seller orders contract buyers according to $\bar{p}_{jt}(1 - \theta_{ijt}) - c_{ijt}$ and trades Y_{ijt} with all but the marginal contract buyer, then the result will be a self-reinforcing solution to (2).

We parameterize heterogeneity in the spot match probability γ_{ijt} as

$$\gamma_{ijt} = \exp(X'_{ijt}\beta_\gamma)/(1 + \exp(X'_{ijt}\beta_\gamma)).$$

Because X_{ijt} includes an indicator for whether a buyer ever signs a contract, the spot match probability will be heterogeneous for contract and spot buyers. This heterogeneity is important because contract buyers may be selected based on their spot match probability.

We specify that bargaining power δ_{ijy} is constant and equal to δ. To explain the fact that contract buyers do not trade in 11% of months, we assume that there exists a contractible variable ϕ_{ijt} such that the value of trade is negative infinity whenever ϕ_{ijt} is negative. We refer to ϕ_{ijt} as the buyer's 'trade type.' This variable does not enter into the spot problem, and its existence only decreases the value of quantity contracts. We specify $\phi_{ijt} = X'_{ijt}\beta_\phi + \epsilon^\phi_{ijt}$, where ϵ^ϕ_{ijt} are i.i.d. and follow a type I extreme value distribution. We assume 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.3 shows that the probability of trade is statistically flat over the average margin for contract buyers.

We specify that the contract buyers in market j and year y are selected according to

$$C_{jy} = \{i : V_{ijy} \geq V^D_{ijt} \text{ with probability } 1\} = \{i : V_{ijy} + X'_{ijy}\beta_D + \sigma_D\epsilon^D_{ijy} \geq 0\}, \text{ where } \epsilon^D_{ijy} \overset{iid}{\sim} TIEV.$$

In this specification, the covariate vector X_{ijy} is the average value of X_{ijt} across months in year y.

5.2 Identification

The model includes twelve parameter vectors that we estimate: $(\beta_\theta, \beta^C_Y, \beta^S_Y, \beta_\gamma, \beta_\phi, \beta_D, \sigma_{\theta\theta}, \sigma^C_{YY}, \sigma^S_{YY}, \sigma_{Y\theta}, \sigma_D, \delta)$. These parameters govern the rebate type distribution, the quantity type distribution for contract and spot buyers, the spot match probability, the contract trade probability, contract selection, and bargaining power. The data are the joint distribution of contract status, rebate, logistics costs, a trade indicator, quantity, and covariate vectors: $(C_{ijy}, R_{ijt}, c_{ijy}, \tau_{ijt}, q_{ijt}, X_{ijt}, X^C_{ijt})$. These data point identify the parameters:

Proposition 2. *The joint distribution of $(C_{ijy}, R_{ijt}, c_{ijy}, \tau_{ijt}, q_{ijt}, X_{ijt}, X^C_{ijt})$ point identifies the model parameters $(\beta_\theta, \beta^C_Y, \beta^S_Y, \beta_\gamma, \beta_\phi, \beta_D, \sigma_{\theta\theta}, \sigma^C_{YY}, \sigma^S_{YY}, \sigma_{Y\theta}, \sigma_D, \delta)$.*

We provide a proof in Online Appendix B.1. The proof does not rely on the assumption that $(\epsilon^\theta_{ijt}, \epsilon^Y_{ijt})$ is normally distributed, which is made to facilitate simulation-based estimation. Variation in the minimum spot price combined with the joint distribution of spot buyer rebates and quantities identify the joint distribution of rebate and quantity types $(\beta_\theta, \beta^S_Y, \sigma_{\theta\theta}, \sigma^S_{YY}, \sigma_{Y\theta})$. This is true because, for all but the marginal spot buyer, rebates and quantities equal rebate type and quantity type. This distribution is truncated based on the minimum spot price. As the minimum spot price approaches the infimum of the support, we observe the full distribution. Next, the way the spot match probability changes with the covariates, β_γ, is identified from the correlation between trade and price in the spot market. The stronger this correlation, the more selective the seller is in choosing spot buyers, meaning the spot match probability is higher. Table OA.3 shows a positive correlation between average spot buyer margin and trade probability, evidence that the seller is somewhat selective in choosing spot buyers.

The distribution of quantity conditional on trade identifies the distribution of quantity type for contract buyers $(\beta_Y^C, \sigma_{YY}^C)$. The probability of trade for contract buyers identifies β_ϕ. The contract value V_{ijy} is a function of the joint type distribution, the spot match probability, and the contract trade probability. The contract selection parameters (β_D, σ_D) are identified by the correlation between V_{ijy}, X_{ijy}, and contract status C_{ijy}. Finally, inversion of (3) combined with identification of the other model parameters identifies bargaining power δ. A key assumption is that the distribution of rebate types is the same for spot and contract buyers. Alternatively, differences in transfers for contract buyers could be driven by different distributions of outside options. We explore this alternative model in Section 7.4.

5.3 Estimation

We estimate the parameters in two steps. We provide further details in Online Appendix B.2.

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: $(\beta_\theta, \beta_Y^S, \sigma_{\theta\theta}, \sigma_{YY}^S, \sigma_{Y\theta}, \beta_\gamma)$. 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:

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

The first five moment vectors match the mean and variance-covariance matrix of rebate and quantity to the

Figure 4: Distribution of spot match probability by contract status

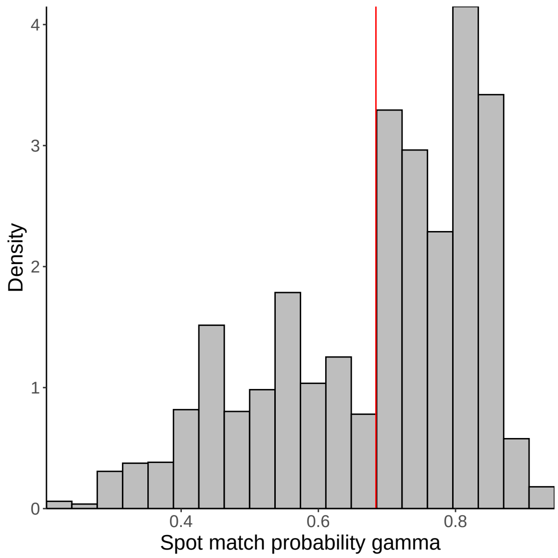

(a) Contract buyers

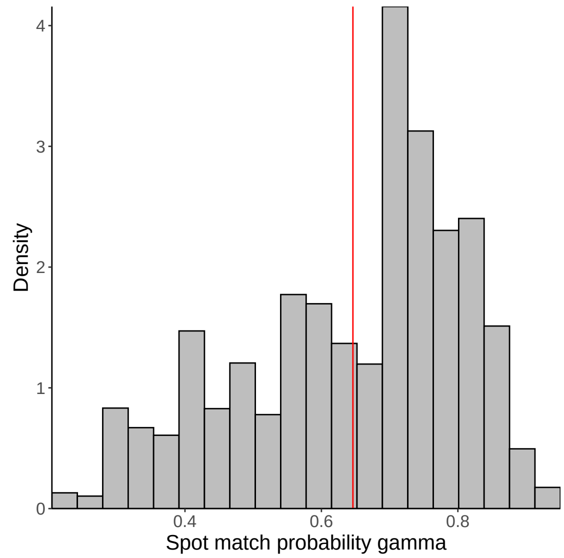

(b) Spot buyers

Note. Each observation is a buyer-market-month. The red bar indicates the mean of each distribution. The mean of the spot match probability for contract buyers is 0.68 and the mean of the spot match probability for spot buyers is 0.65.

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: $(\beta_Y^C, \beta_\phi, \beta_D, \sigma_{YY}^C, \sigma_D, \delta)$. To estimate these parameters, we run linear and logistic 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. Then, using the parameters that have already been estimated, we simulate the contract value V_{ijy} and run a logistic regression of contract status on V_{ijy} and X_{ijy} to estimate (β_D, σ_D). To estimate δ, we first simulate the terms in (3) using the parameters that have already been estimated. We then estimate δ based on equation (3) and the Nash-in-Nash constraints.

5.4 Estimates

Figure 4 plots the distribution of estimated spot match probabilities by contract status. We estimate that average spot match probability γ_{ijt} equals 0.66 with standard error 0.06. This means that the seller expects to come into contact with 66% of the spot buyers in each month. For comparison, the average probability

Table 4: Parameter estimates

Parameter	Spot match prob. β_γ	Contract selection β_D	Rebate type β_θ	Quantity type (spot) β_Y^S	Quantity type (contr.) β_Y^C	Trade type β_ϕ
Ever contract	0.63	1.05	0.03	-0.22		-0.13
	(1.80)	(6.39)	(0.02)	(0.04)		(0.17)
Avg. Quantity >1,000 tons	-0.02	1.14	0.03	1.11		1.18
	(0.35)	(3.08)	(0.008)	(0.06)		(0.15)
Prioritized	0.56	1.35	-0.07	-0.22		0.07
	(0.26)	(2.66)	(0.008)	(0.03)		(0.15)
Logistics Costs	-0.19	-0.96	-0.06	-0.04	0.20	-1.20
	(0.97)	(3.35)	(0.02)	(0.07)	(0.29)	(0.33)
Market Price	0.08	-0.19	0.03	0.04	-0.06	0.19
	(0.08)	(2.54)	(0.002)	(0.01)	(0.03)	(0.05)
Time Trend (months)	-0.01	0.01	-0.002	0.0004	0.003	-0.008
	(0.008)	(3.26)	(0.0002)	(0.001)	(0.004)	(0.003)
Market FE	Yes	Yes	Yes	Yes		Yes
Buyer-market-year FE					Yes	
Mean spot trade probability γ_{ijt}		0.66 (0.06)				
Buyer bargaining power δ		0.29 (0.003)				
Variance in contract selection σ_D		0.86 (33.61)				

Notes. Analytic standard errors are in parentheses for all estimates except for the contract selection parameters and bargaining power, where standard errors are calculated by a parametric bootstrap. Bootstrapped standard errors are in parentheses. Quantity type is in thousands of tons. Logistics costs is in hundreds of euros. Market price is in hundreds of euros.

of trade among spot buyers is 55%. Therefore, the seller is somewhat selective about the spot buyers with whom it trades. However, trading frictions are still sizable because we estimate γ_{ijt} less than one. There is heterogeneity in the estimated spot match probability by contract status. For contract buyers, the average spot match probability is 0.68 and the spot buyers the average spot match probability is 0.65.

Table 4 presents the estimates of the other model parameters. We do not find statistically detectable selection of contract buyers based on observables. However, some of the point estimates are relatively large and have expected signs, as they suggest that large or prioritized buyers with low logistics costs are the most likely to sign a contract. We estimate that buyer bargaining power δ is 0.29. Institutional details suggest that the seller should have most of the relative bargaining power because the seller is larger than most buyers. Furthermore, the buyer bargaining power δ 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}. Therefore, δ may be small even if buyers have some bargaining power that enters into the rebate type. We evaluate the robustness of our results to an alternative contract buyer pricing model where the seller makes take-it-or-leave-it offers in Section 7.4.[20]

[20]This alternative analysis provides a conservative estimate for the value of quantity contracts because it loads the entire difference in contract buyer rebates onto contract buyer outside options, depressing the assurance value. We do not consider an analysis where the buyers make take-it-or-leave-it offers because doing so would over-estimate the value of quantity contracts. Furthermore, such a model would not explain variation in contract buyer rebates within a market-month.

Figure 5: The value of quantity contracts relative to profits

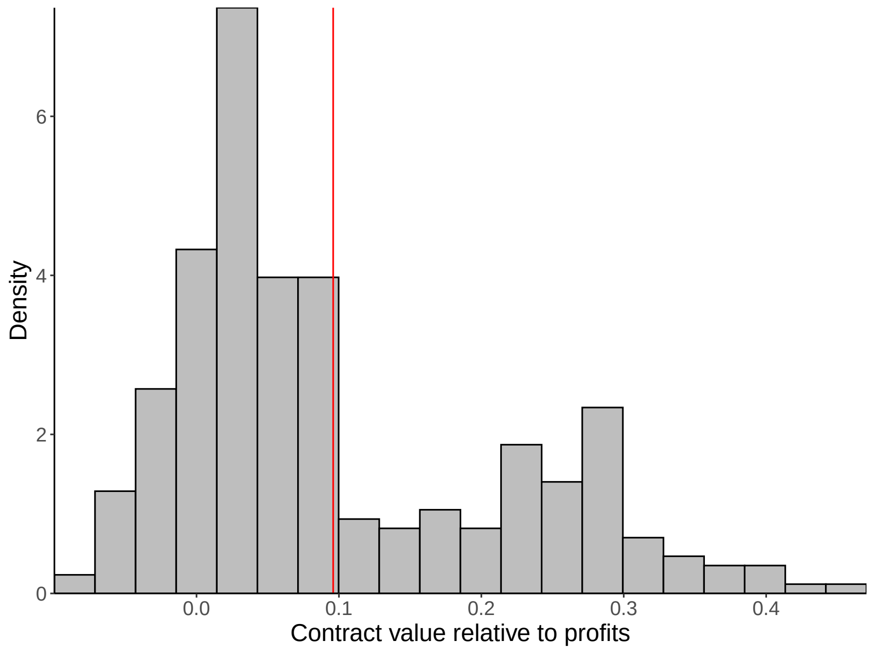

Note. Each observation is a contract buyer within a market and year. The red line at 0.096 corresponds to the mean of the distribution. The contract value equals the difference in buyer and seller payoffs when a buyer trades with a contract and when a buyer trades on the spot. Profits are the average of the buyer's price type (market price times one minus rebate type) minus logistics costs times quantity. We measure profits with price type instead of price because negotiated prices are a function of the contract value.

6 The Value of Quantity Contracts

Equipped with the model estimates, we estimate the value of each contract. We calculate the difference in expected payoffs when a contract buyer signs a contract versus when that buyer trades on the spot, holding fixed the set of other contracts. To interpret this difference, we divide by profits—the average margin times quantity over the year. Figure 5 plots the distribution of the value of quantity contracts relative to profits for both contract buyers and spot buyers. For contract buyers, the mean value of a quantity contract is 10% of profits and there is substantial heterogeneity. The ability write quantity contracts increases total profits by 29%. The value is positive for 82% of contracts. The value may be negative for contract buyers due to unfavorable draws of contractible variables X^C_{ijt} over the course of the year. For spot buyers, the mean value of a contract is 6% of profits and 37% have a negative counterfactual value. The value may be positive for spot buyers because imperfect enforcement prevents the spot buyer from signing a valuable contract.

Next, we decompose the value of quantity contracts into an assurance value and a lost option value in Table 5. To calculate the assurance value, we simulate the expected difference in payoffs if a contract buyer trades the contracted quantity versus if the contract buyer does not trade at all. We calculate the lost option value using the decomposition of the contract value into a weighted average of the assurance value and the

Table 5: Decomposition into assurance value and lost option value

	N	Mean	SD	p25	Median	p75
Contract value	300	0.096	0.117	0.018	0.051	0.183
Assurance value	300	0.411	0.178	0.285	0.386	0.531
Lost option value	300	0.089	0.095	0.032	0.068	0.104
Weight on assurance value	300	0.316	0.152	0.182	0.278	0.442

Note. Each observation is a contract buyer within a market and year. The assurance value is the difference in expected payoffs when a contract buyer trades versus when a contract buyer does not trade. The lost option value is the difference in expected payoffs when a contract buyer matches with the seller on the spot versus when the contract buyer trades the contracted quantity. The weight on the assurance value is one minus the spot match probability.

lost option value. The mean assurance value is 41% of profits. This means that the seller can recoup more than half of the lost profits on the spot market if a contract buyer were to not trade in some period. The mean lost option value is 9% of profits. Consistent with the stylized fact that there is significant ex-post deviation between the minimum spot and contract margins, payoffs would be significantly higher if parties could contract on spot market realizations. The average contract value would be 15% of profits absent the lost option value.

Figure 6: Comparative statics of contract value

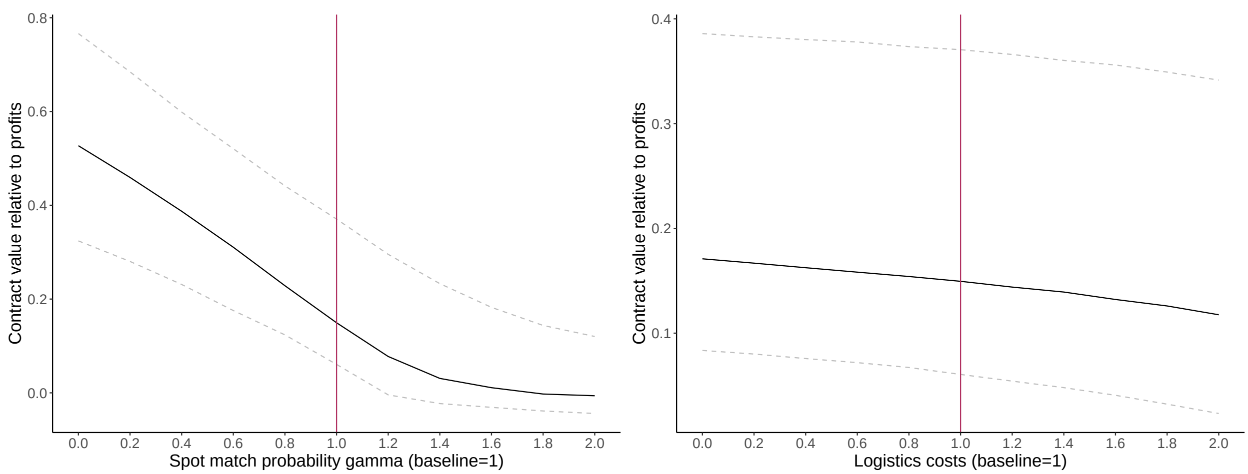

(a) Spot match probability

(b) Logistics costs

Note. The solid line indicates the median of the distribution of contract value relative to profits and the the dashed lines correspond to the interquartile range. In the left figure, the spot match probability of each contract buyer is multiplied by the number on the x-axis, and truncated to [0, 1]. In the right figure, the logistics costs for each contract buyer is multiplied by the number on the x-axis. The red line indicates the baseline specification.

Figure 6 presents the results of two comparative statics. The left panel plots how the distribution of the value of quantity contracts changes with the probability of matching on the spot γ_{ijt}. The value decreases in this probability. A 10% increase in the probability of matching on the spot is associated with a reduction in the mean value of a quantity contract by 26%. As the spot match probability approaches one, the seller can

Figure 7: Trading frictions and contracting behavior

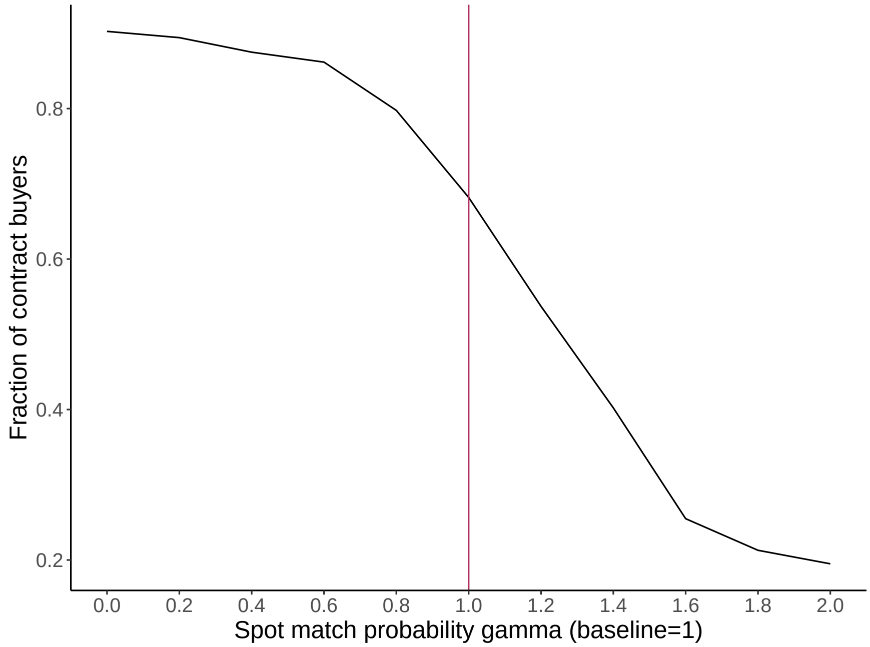

Note. This figure plot the fraction of all buyers that sign a contract as the spot match probability is multiplied by the value on the x-axis for all buyers and truncated to $[0, 1]$. The red line indicates the baseline specification. The equilibrium plotted here is the closest equilibrium to the observed set of contract buyers.

trade with any of the spot buyers, so it has maximal choice. If the seller signs a contract, then it restricted its choice set. As a result, the seller prefers to trade on the spot. The right panel of Figure 6 plots how the distribution of the value of quantity contracts changes with contract buyer logistics costs. A 10% increase in logistics costs causes a reduction in the mean value of a contract by 3%.

The next exercise considers how contracting behavior changes with the spot match probability. Figure 7 plots the fraction of buyers with a contract as the spot match probability increases. In this exercise, we increase the spot match probability for all buyers and use the contract selection model to predict the set of equilibrium contracts. The fraction of buyers with a contract declines with the spot match probabilities because contract value declines with spot match probability. In the baseline, 68% of buyers sign contracts. As the spot match probabilities approach one for all buyers, the fraction of buyers signing contracts approaches 19%. Figure 8 shows how profits change with the spot match probability. As the spot match probability increases, total contracted profits fall, but the rise in spot profits more than compensates causing total profits to increase. If spot match probabilities doubled, total profits would increase by 3.4% and contracted profits would fall by 82%. This counterfactual offers an estimate for the gains the seller could achieve from an investment in contacting spot buyers earlier and more often, holding fixed the behavior of other sellers in the industry.

Figure 8: Profits by contract status over spot match probability

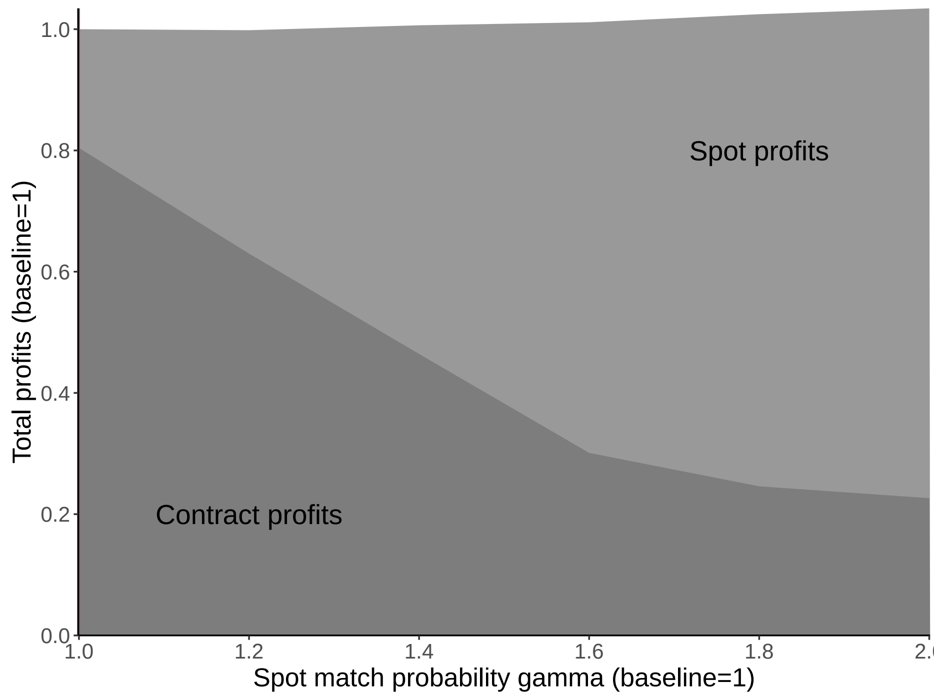

Note. This figure presents profits (equal to margin times quantity) over contract status for varying spot match probabilities. The spot match probability is multiplied by the value on the x-axis for all buyers and truncated to $[0, 1]$. The contract selection model determines the set of contract buyers as the closest equilibrium to the observed allocation.

7 Extensions and Robustness

7.1 Quantity Contracts and Buyer Reliability

A secondary function of quantity contracts may be to induce buyer reliability. In decentralized markets, risk neutral sellers prefer buyers to purchase stable quantities. This preference arises endogenously from the structure of the spot market. In the model presented in Section 3, expected profits are concave with respect to spot quantity Q_{jt}^S. This is true because the marginal spot payoff is decreasing in spot quantity. As a consequence, the risk neutral seller has an endogenous preference for stable contract quantities. Quantity contracts can commit the buyer to trade stable quantities if buyers have access to storage technology.

Empirically, contract buyers are more reliable than spot buyers. Conditional on trade, the coefficient of variation of contract buyer quantity is 9% smaller than that of spot buyer quantity. Table OA.4 shows this difference is robust to the inclusion of buyer fixed effects. To estimate how the value of quantity contracts changes if they increase buyer reliability, we simulate a counterfactual where contracts reduce the coefficient of variation of quantity while keeping fixed mean quantity conditional on trade. Figure OA.4 presents the results. If quantity contracts reduce the coefficient of variation of quantity by 9%, then the average value of a quantity contract would 27% of profits.

7.2 Quantity Contracts and Price Dispersion

The use of quantity contracts can affect equilibrium price dispersion. We measure price dispersion as the expected difference between the maximum and minimum price within a market and month. Quantity contracts affect price dispersion through two channels. First, quantity contracts insure the seller against having to trade with undesirable spot buyers in the contingency where the contract buyer would not match with the seller as a spot buyer. This effect is magnified if the probability of the seller matching with the contract buyer on the spot is small, or if the difference between the contract buyer's valuation and the worst spot buyer's valuation is large. Quantity contracts decrease price dispersion through this channel. Second, quantity contracts prevent the seller from ex-post optimizing the quantity it trades with the contract buyer. If the seller experiences a bad draw from the spot market, then the seller is unable to increase the quantity it trades with the contract buyer. As a result, the minimum price it receives is lower than the price it could get if it matched with the contract buyer on the spot in that same spot market. Therefore, quantity contracts increase price dispersion through this channel.

The above two channels operate in opposite directions, so the net effect of quantity contracts on price dispersion is ambiguous. Figure OA.5 plots the histogram of how quantity contracts change price dispersion. For each contract buyer i, we simulate price dispersion if buyer i traded on the spot. On average, quantity contracts reduce price dispersion by 3%, but in some markets reduce price dispersion by as much as 27%.

7.3 Inventory

We consider an extension where the seller utilizes inventory to smooth total sales. The ability to utilize inventory may diminish the value of quantity contracts because it decreases the need for quantity insurance against undesirable spot market outcomes. In the extension, we estimate that the mean value of a quantity contract is 18% of profits when the seller can use inventory, which is even higher than the baseline estimate, so inventory does not meaningfully diminish the value of quantity contracts. Part of this difference could be due to sampling variation or model misspecification. We provide details in Online Appendix C.1.

To incorporate inventory into the model, we suppose that the seller chooses quantity to trade off between spot buyer margins and the marginal net value of inventory. We identify the marginal net value of inventory using the negative correlation between the minimum spot margins and inventory levels. Though this correlation is statistically significant, it is economically small, explaining why the incorporation of inventory does not dampen the main results. One explanation for this small correlation may be that the logistical costs of inventory management are high in this setting.

7.4 Alternative Contract Buyer Pricing Model

We consider an alternative contract buyer pricing model where the seller makes take-it-or-leave-it offers. This model allows us to identify the contract buyer rebate type from the observed rebates. We allow for rich heterogeneity in the contract buyer rebate type distribution:

$$\theta_{ijt} = C_{ijt} X_{ijt}^{C\prime} \beta_{\theta}^{C} + (1 - C_{ijt}) X_{ijt}^{\prime} \beta_{\theta}^{S} + \epsilon_{ijt}^{\theta}.$$

We estimate the new parameter vector β_{θ}^{C} using ordinary least squares. We present the results in Table OA.5.

Under this alternative pricing model, we estimate that the average value of a quantity contract is 10% of profits, a difference of less than one percentage point from the baseline estimate. Figure OA.6 includes the histogram of quantity contract values. The value is positive for 95% of contracts, compared with 82% under the baseline model.

8 Discussion

8.1 General Relevance

Many industries rely on quantity contracts or similar arrangements to organize trade. When might an industry use quantity contracts to mitigate the costs of trading frictions, as in the pulp industry? First, our analysis suggests that logistical frictions are of primary importance. Global industries where buyers and sellers are separated by large geographic distances might face large trading frictions and turn to quantity contracts. Second, quantity contracts serve a similar purpose as inventory in the pulp industry. Therefore, industries where inventory is costly due to spoilage or bulk might rely on quantity contracts. Third, quantity contracts in the pulp industry provide assurance that the most valuable trades will occur. Industries where some matches are more valuable than others, perhaps due to product differentiation or geographic location, might turn to quantity contracts. These three features are salient in many firm-to-firm industries, suggesting that our analysis of quantity contracts is broadly relevant.

8.2 Quantity Contracts and Intermediation

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 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 [2022] 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, 2022], contracts may have longer horizon. Additionally, when it is costly to select suppliers, firms may use contracts with longer horizon [MacKay, 2022].

9 Conclusion

In many industries, buyers and sellers rely on quantity contracts to organize trade. We show that quantity contracts are valuable in decentralized markets because they reduce the costs of trading frictions. Quantity contracts provide quantity assurance because trading frictions may prevent a buyer and seller from matching on the spot. However, quantity contracts inhibit sellers from optimally allocating their production across buyers after market conditions realize. We develop an empirical model of quantity contracts and quantify these forces in the pulp and paper industry. We find that the mean value of a quantity contract is 10% of profits and the lost option value is sizable.

References

Jason Allen, Robert Clark, and Jean-François Houde. Search frictions and market power in negotiated-price markets. *Journal of Political Economy*, 127(4):1550–1598, 2019.

Treb Allen. Information frictions in trade. *Econometrica*, 82(6):2041–2083, 2014.

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

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

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

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

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

Ambarish Chandra and Mariano Tappata. Consumer search and dynamic price dispersion: an application to gasoline markets. *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.

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

José Ignacio Cuesta, Carlos Noton, and Benjamin Vatter. Vertical integration between hospitals and insurers. 2019. Working paper.

Maryam Farboodi, Gregor Jarosch, and Guido Menzio. Intermediation as rent extraction. 2019. Working paper.

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

Alessandro Gavazza and Alessandro Lizzeri. Frictions in product markets. In *Handbook of Industrial Organization*, volume 4, pages 433–484. Elsevier, 2021.

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

Adam Harris and Thi Mai Anh Nguyen. Long-term relationships in the US truckload freight industry. 2021. Working paper.

Adam Harris and Thi Mai Anh Nguyen. Long-term relationships and the spot market: Evidence from US trucking. 2023.

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. Price adjustment in long-term contracts: The case of coal. *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, 2013.

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. Relational contracts and development. *Annual Review of Economics*, 14:337–362, 2022.

Rocco Macchiavello and Josepa Miquel-Florensa. Vertical integration and relational contracts: Evidence from the Costa Rica coffee chain. 2018. Working paper.

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.

Alexander MacKay. Contract duration and the costs of market transactions. *American Economic Journal: Microeconomics*, 14(3):164–212, 2022.

Guillermo Marshall. Search and wholesale price discrimination. *The RAND Journal of Economics*, 51(2):346–374, 2020.

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. *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.

Claudia Robles-Garcia. Competition and incentives in mortgage markets: The role of brokers. 2020. Working paper.

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. *Journal of Political Economy*, 130(2), 2022.

Victor Stango and Jonathan Zinman. Borrowing high versus borrowing higher: Price dispersion and shopping behavior in the US credit card market. *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.

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

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. 2022. Working paper.

Nahim Bin Zahur. Market power and spatial price discrimination in the liquefied natural gas industry. 2023. Working paper.

Appendix

For Online Publication

Pulp Friction: The Value of Quantity Contracts in Decentralized Markets

Olivier Darmouni Simon Essig Aberg Juha Tolvanen

Online Appendix A Additional Exhibits

Figure OA.1: Timeseries of Gross Price by Index Status

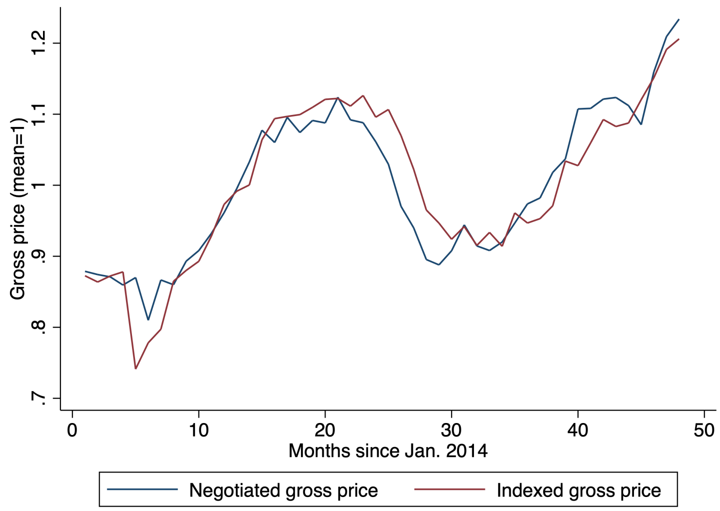

Notes. The sample is restricted to those observations where we observe an indicator of whether the contract indexed price or negotiated price. Both series plotted relative to their means.

Table OA.1: Market price elasticity of quantity sold, production, and inventory

	Quantity sold (log)	Production (log)	Inventory (log)	Quantity sold (log)	Production (log)	Inventory (log)
Market price (log)	-0.092	-0.057	-0.216	-0.105	0.076	-0.088
	(0.117)	(0.180)	(0.143)	(0.094)	(0.149)	(0.175)
Market FEs	Yes	Yes	Yes			
Market-year FEs				Yes	Yes	Yes
Observations	288	288	288	288	288	288
R^2	0.898	0.694	0.927	0.925	0.729	0.953

Note. The unit of observation is a market-month. Production and inventory are reported at the fiber-mill-month level in the raw data, so these variables are aggregated to the market-month level by taking a quantity-weighted average within a market and month.

Table OA.2: Price elasticity of quantity demanded

	Price (log)	Quantity (log)	Price (log)	Quantity (log)
Logistics costs (log)	0.031		0.087	
	(0.004)		(0.022)	
Price (log)		-17.5		0.362
		(3.31)		(0.810)
Model	OLS	2SLS	OLS	2SLS
Market-year fixed effects	Yes	Yes		
Buyer-market-year fixed effects			Yes	Yes
Observations	6,489	6,489	6,489	6,489
R^2	0.515		0.779	
Robust F-statistic		69.3		15.5

Note. Each observation is a buyer-market-month conditional on trade. The price measure is price after rebates. In the second and fourth column, logged logistics costs instrument for price.

Figure OA.2: Fraction trading contract among buyers that switch status

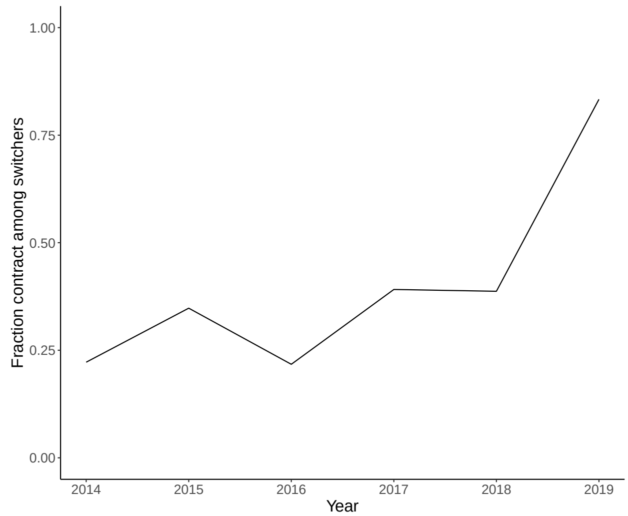

Note. This figure plots the fraction of buyer-markets that trade using contracts, among buyer-markets that switch status between contract and spot at least once in the sample period.

Figure OA.3: Difference in trade probability after moving from contract to spot

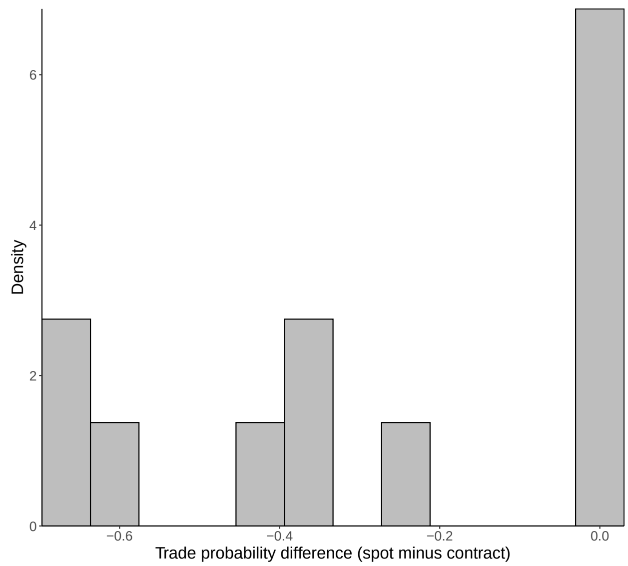

Note. Each observation is a buyer-market-year. The sample includes the twelve observations that were classified as contract in year $t-1$ and spot in year t. The average trade probability while in contract in year $t-1$ is 0.84 and the average trade probability while in spot in year t is 0.57.

Table OA.3: Patterns in Probability of Trade

	Pr(trade) (1)	Pr(trade) (2)	Pr(trade) (3)	Pr(trade) (4)
Contract Buyer	0.35	0.28	0.68	0.54
	(0.02)	(0.02)	(0.11)	(0.13)
Spot buyer × Average margin			0.45	1.08
			(0.14)	(0.30)
Contract buyer × Average margin			-0.20	0.61
			(0.15)	(0.34)
Capacity (Millions Tons)		0.36		0.39
		(0.22)		(0.22)
Priority		0.08		0.06
		(0.02)		(0.02)
Logistics Costs		-0.11		-0.06
		(0.06)		(0.06)
Constant	0.55	0.55	0.32	0.09
	(0.01)	(0.07)	(0.07)	(0.15)
Market-year FE	NO	YES	NO	YES
Observations	797	578	797	578
R-squared	0.27	0.40	0.28	0.41

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.

Table OA.4: Contract buyer and spot buyer coefficient of variation in quantity

	C.V. quantity	C.V. quantity	C.V. quantity (cond. on trade)	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)
Market-year FE	NO	YES	NO	YES
Observations	797	574	797	572
R-squared	0.25	0.37	0.03	0.10

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.4: Comparative statics of contract value with C.V. of spot quantity

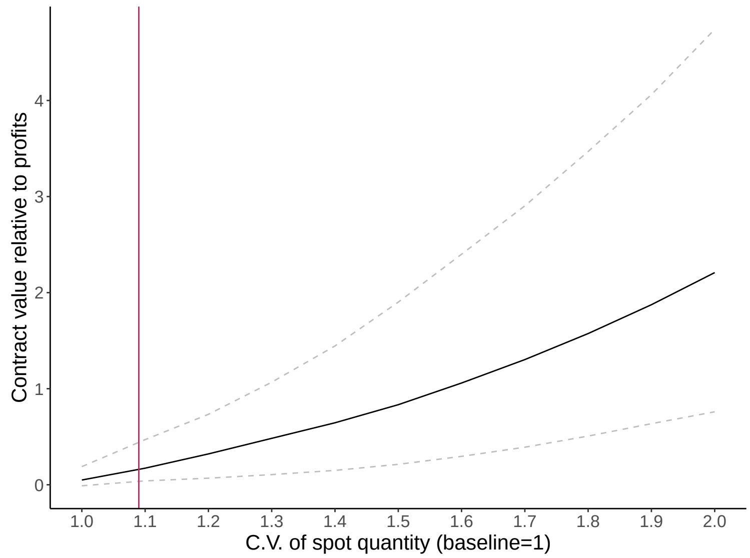

Notes. The solid line indicates the median of the distribution of contract value relative to profits and the the dashed lines correspond to the interquartile range. The coefficient of variation equals the standard devation divided by the mean. The coefficient of variation of quantity conditional on trade for spot trade is multiplied by the value on the x-axis, keeping fixed the mean. Quantity conditional on trade is the maximum of a normal random variable and zero. Therefore, to make the requisite adjustment to the coefficient of variation keeping fixed the mean of this random variable, we numerically search for the adjusted mean and standard deviation of the latent normal random variable. We then calculate the value of a quantity contract when buyers have the adjusted mean and standard deviation when trading on the spot.

Figure OA.5: Distribution of Price Dispersion without Contract

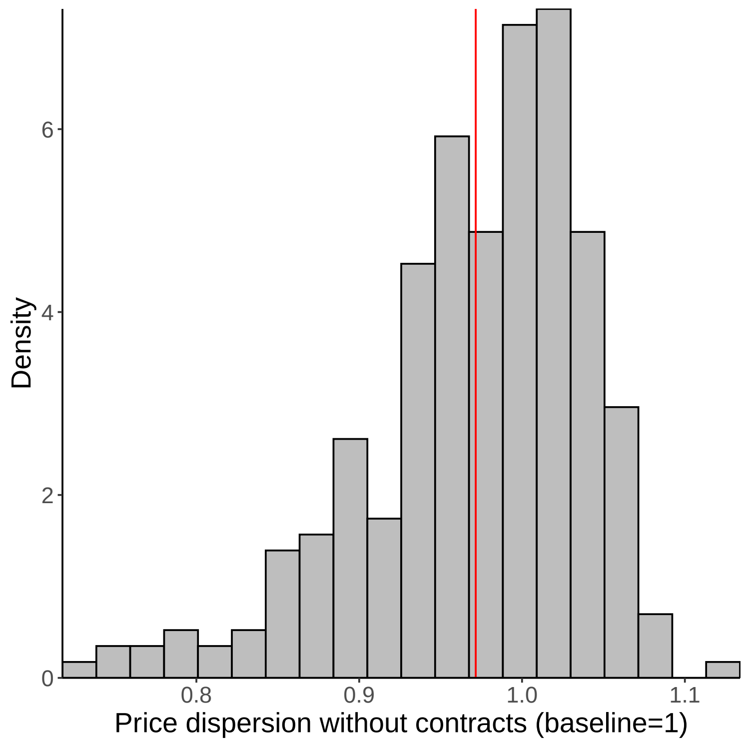

Notes. Each observation corresponds to a market and month. Price dispersion equals the expected difference between the maximum and minimum price the seller receives. Price dispersion when there are no contracts is measured relative to the observed model-predicted price dispersion. The mean of 0.97 is indicated in red.

Table OA.5: Parameter estimates, alternative pricing model

Parameter	Spot match prob. β_γ	Rebate type (spot) β_θ^S	Rebate type (contr.) β_θ^C	Quantity type (spot) β_Y^S	Quantity type (contr.) β_Y^C	Trade type β_ϕ
Ever contract	0.63		0.03	-0.22		-0.13
Avg. Quantity >1,000 tons	-0.02		0.03	1.11		1.18
Prioritized	0.56		-0.07			0.07
Logistics Costs	-0.19	-0.07	-0.06	-0.04	1.49	-1.20
Market Price	0.08	0.01	0.03	0.04	-0.09	0.19
Time Trend (months)	-0.01	-0.002	-0.002	0.0004	0.006	-0.008
Market FE	Yes		Yes	Yes		Yes
Buyer-market-year FE		Yes			Yes	

Notes. Quantity type is in thousands of tons. Logistics costs is in hundreds of euros. Market price is in hundreds of euros.

Figure OA.6: Value of quantity contracts, alternative pricing model

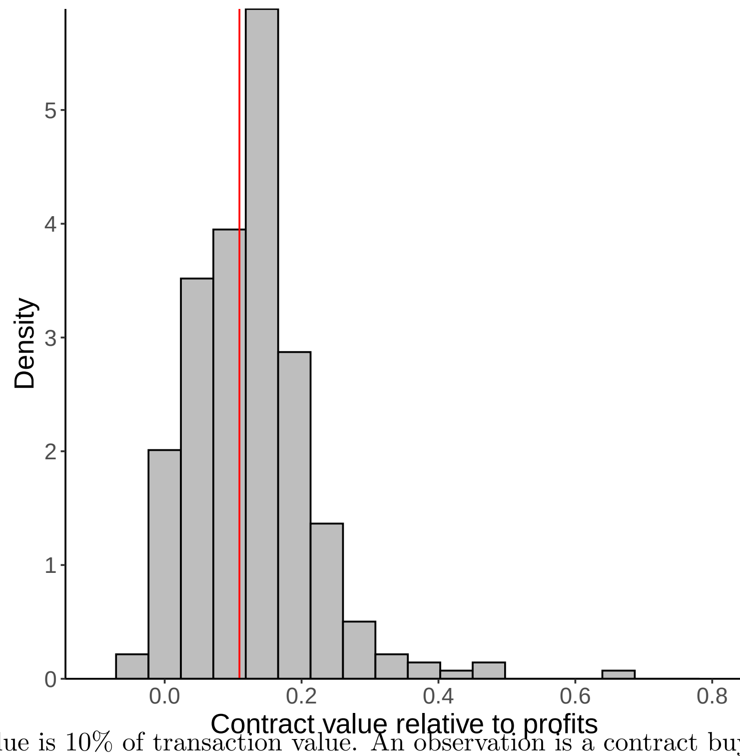

Notes. The mean value is 10% of transaction value. An observation is a contract buyer within a market and year.

Online Appendix B Details on Identification and Estimation

Online Appendix B.1 Identification

In this section, we prove that the model parameters are identified.

Proposition 2. *The joint distribution of* $(C_{ijy}, R_{ijt}, c_{ijy}, \tau_{ijt}, q_{ijt}, X_{ijt}, X^C_{ijt})$*. point identifies the model parameters* $(\beta_\theta, \beta^C_Y, \beta^S_Y, \beta_\gamma, \beta_\phi, \beta_D, \sigma_{\theta\theta}, \sigma^C_{YY}, \sigma^S_{YY}, \sigma_{Y\theta}, \sigma_D, \delta)$*.*

Proof. The joint distribution of rebate type and quantity type for spot buyers $(\beta_\theta, \beta^S_Y, \sigma_{\theta\theta}, \sigma^S_{YY}, \sigma_{Y\theta})$*.* The joint distribution of rebate type and quantity type is identified from the distribution of prices and quantities conditional on trade in the spot market. The analysis of the seller's spot market problem implies that

$$\lim_{|N_{jt}|\to\infty} P(q_{ijt} = Y_{ijt} | \tau_{ijt} = 1, X_{ijt}) = 1.$$

Let $\tilde{p}_{ijt} = \bar{p}_{jt}(1-\theta_{ijt}) - c_{ijt}$ denote the margin type of buyer i in market j and month t. When trade occurs, $p_{ijt} = \tilde{p}_{ijt}$. Let $\underline{p}_{jt}$ denote the minimum margin in market j and month t. Conditional on matching with the seller, trade occurs if and only if $\tilde{p}_{ijt} \geq \underline{p}_{jt}$. Therefore,

$$(p_{ijt}, q_{ijt}) | \underline{p}_{jt}, \tau_{ijt} = 1, X_{ijt} \sim (\tilde{p}_{ijt}, Y_{ijt}) | \tilde{p}_{ijt} \geq \underline{p}_{jt}, X_{ijt}.$$

Taking the limit as $\underline{p}_{jt} \to \inf \operatorname{Supp} \tilde{p}_{ijt}$, we have

$$(p_{ijt}, q_{ijt}) | \underline{p}_{jt}, \tau_{ijt} = 1, X_{ijt} \rightsquigarrow (\tilde{p}_{ijt}, Y_{ijt}) | X_{ijt},$$

where $\rightsquigarrow$ indicates the limit of the conditional distribution. Consequently the joint distribution of $\tilde{p}_{ijt}$ and Y_{ijt} is identified, implying the joint distribution of $\theta_{ijt} = 1 - (\tilde{p}_{ijt} + c_{ijt})/\bar{p}_{jt}$ and Y_{ijt} is identified.

Spot match probability β_γ*.* Spot match probability parameters β_γ are identified from the correlation between trade and price in the spot market. We have

$$\begin{aligned}
&\mathbb{E}[(p_{ijt} - c_{ijt})\tau_{ijt} | X_{ijt}] = P(\tau_{ijt} = 1 | X_{ijt}) \mathbb{E}[p_{ijt} - c_{ijt} | \tau_{ijt} = 1, X_{ijt}] \\
&= P(m_{ijt} = 1 | X_{ijt}) P(\tau_{ijt} = 1 | m_{ijt} = 1, X_{ijt}) \mathbb{E}[p_{ijt} - c_{ijt} | \tau_{ijt} = 1, X_{ijt}] \\
&= \frac{\exp(X'_{ijt}\beta_\gamma)}{1 + \exp(X'_{ijt}\beta_\gamma))} \mathbb{E}[P(\tau_{ijt} = 1 | \underline{p}_{jt}, m_{ijt} = 1)] \mathbb{E}[p_{ijt} - c_{ijt} | \tau_{ijt} = 1, X_{ijt}]
\end{aligned}$$

$$= \frac{\exp(X'_{ijt}\beta_\gamma)}{1+\exp(X'_{ijt}\beta_\gamma))}\mathbb{E}[P(\tilde{p}_{ijt} \geq \underline{p}_{jt}|\underline{p}_{jt}, X_{ijt})]\mathbb{E}[p_{ijt} - c_{ijt}|\tau_{ijt} = 1, X_{ijt}]$$

$$\Rightarrow \frac{\exp(X'_{ijt}\beta_\gamma)}{1+\exp(X'_{ijt}\beta_\gamma))} = \frac{\mathbb{E}[(p_{ijt} - c_{ijt})\tau_{ijt}|X_{ijt}]}{\mathbb{E}[P(\tilde{p}_{ijt} \geq \underline{p}_{jt}|\underline{p}_{jt}, X_{ijt})]\mathbb{E}[p_{ijt} - c_{ijt}|\tau_{ijt} = 1, X_{ijt}]}. \tag{6}$$

The denominator of the right hand side of (6) is identified because the marginal distribution of $\tilde{p}_{ijt}$ is identified. Therefore, β_γ is identified via logit inversion.

The marginal distribution of quantity type for contract buyers (β^C_Y, σ^C_{YY}). The analysis of the contract problem implies that

$$\lim_{|C_{jy}|\to\infty} P(q_{ijt} = Y_{ijt}|\tau_{ijt} = 1, X^C_{ijt}) = 1.$$

Hence,

$$q_{ijt}|\tau_{ijt} = 1, X^C_{ijt} \sim Y_{ijt}|\tau_{ijt} = 1, X^C_{ijt},$$

meaning the marginal distribution of quantity type for contract buyers is identified.

Contract buyer trade type β_ϕ. The distribution of ϵ^ϕ_{ijt} implies that for contract buyers

$$\mathbb{E}[\tau_{ijt}|X_{ijt}] = \frac{\exp(X'_{ijt}\beta_\phi)}{1+\exp(X'_{ijt}\beta_\phi)},$$

so β_ϕ is identified via logit inversion.

Contract buyer bargaining power δ_{ijy}. Contract buyer bargaining power is identified by inverting equation (3). All remaining terms in (3) are functions of the previously identified parameters, so δ_{ijy} is identified.

Contract selection parameters (β_D, σ_D). The value of a contract V_{ijy} is a function of the previously identified parameters. The contract selection model implies

$$\mathbb{E}[C_{ijy}|V_{ijy}, X_{ijy}] = \frac{\exp((V_{ijy} + X'_{ijy}\beta_t)/\sigma_D)}{1+\exp((V_{ijy} + X'_{ijy}\beta_t)/\sigma_D)}.$$

Therefore, the prameters of the contract selection model are identified via logit inversion.

□

Online Appendix B.2 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: $(\beta_\theta, \beta_Y^S, \sigma_{\theta\theta}, \sigma_{YY}^S, \sigma_{Y\theta}, \beta_\gamma)$. We use two-stage generalized method of moments. There are six sets of moments:

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

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 is selected at random. The probability that $i \in S$ is in N is given by $\gamma_i = \exp(X_i'\beta_\gamma)/(1 + \exp(X_i'\beta_\gamma))$ and buyers are selected independently.

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:

$$\begin{aligned}
&\mathbb{E}[\tau_i|\theta_i, Y_i, i \in N, X_i] \\
&= Pr(\tilde{Y}_i \leq Q^S|\theta_i, Y_i, i \in N, X_i) \quad (7)\\
&= Pr\left(\sum_{i' \in N: p_{i'} > p_i} Y_{i'} + Y_i \leq Q^S \middle| \theta_i, Y_i, i \in N, X_i\right) \quad (8)\\
&= Pr\left(\frac{1}{|i' \in N : p_{i'} > p_i|} \sum_{i' \in N: p_{i'} > p_i} Y_{i'} \leq \frac{1}{|i' \in N : p_{i'} > p_i|}(Q^S - Y_i) \middle| \theta_i, Y_i, i \in N, X_i\right). \quad (9)
\end{aligned}$$

Equation (7) follows by step five of the algorithm that determines trade. Equation (8) follows by the definition of $\tilde{Y}_i$. Equation (9) 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 $|N|$ 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:

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

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 $\gamma_i = \exp(X_i'\beta_\gamma)/(1 + \exp(X_i'\beta_\gamma))$ that is independent of the buyer's type. Therefore,

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

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: $(\beta_Y^C, \beta_\phi, \beta_D, \sigma_{YY}^C, \sigma_D, \delta_{ijy})$. 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 equation (3) and the Nash-in-Nash constraints. The model implies:

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

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. Based on the Nash-in-Nash constraints, we constrain these estimates to $[0, 1]$. This generates a contract-specific bargaining power. We take the mean of these. Fourth, to estimate the contract selection parameters (β_D, σ_D), we run a logistic regression of contract status C_{ijy} on the value of the contract V_{ijy} and X_{ijy}, where

$$\begin{aligned}V_{ijy} =&\mathbb{E}[\text{BUYER OUTSIDE OPTION}] + \mathbb{E}[\text{SELLER SPOT PROFITS, } i \text{ IN CONTRACT}]\\&-\mathbb{E}[\text{SELLER SPOT PROFITS, } i \text{ IN SPOT}].\end{aligned}$$

Online Appendix C   Details on Extensions and Robustness

Online Appendix C.1   Extension with Inventory

In this section, we describe an extension where the seller can inventory pulp instead of facing a hard constraint on total sales.

The base model assumes that total sales are inelastic to market conditions. In the extension, we retain the assumption that production is inelastic to market conditions, but allow the seller to choose inventory in response to market conditions. Figure OA.7 presents a scatter-plot of inventory versus the minimum spot price that the seller can find in a particular market and month. The slope is significantly negative, consistent with a model where the seller faces a diminishing marginal value of inventory. Though the slope is statistically significant, the magnitude is economically small. One standard deviation increase in the minimum spot price predicts a reduction in inventory by less than 4% of a standard deviation.

Figure OA.7: Inventory Choice Responds to Spot Market Outcomes

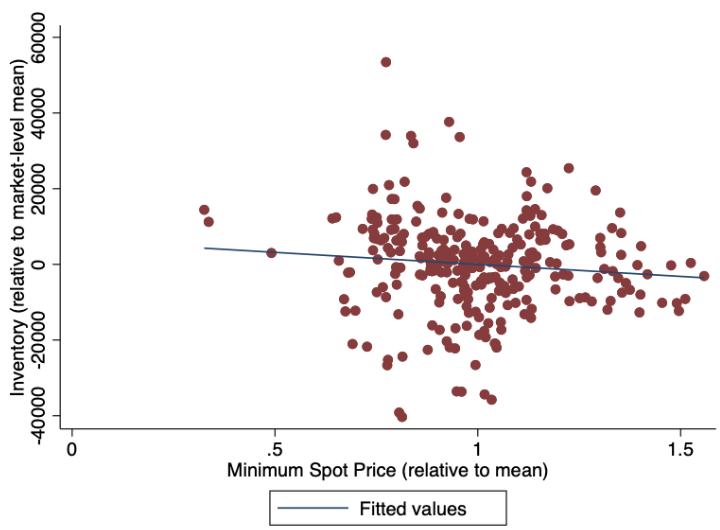

Notes. Observations are at the month-region-fiber level. We observe inventory at the month-fiber-mill level. To construct values at the month-region-fiber level, we consider the average flow of pulp from each mill to each region over the entire period, and take appropriately-weighted averages. Inventory is demeaned relative to region-fiber averages. The price measure is the minimum mill gate price among spot buyers.

The model with inventory builds directly on the base model we develop in Section 3. We suppose that the seller chooses inventory I_{jt} in market j and month t satisfying the accounting identity

$$I_{jt} = I_{jt-1} + Q_{jt} - \sum_i q_{ijt}.$$

Given this identity, the seller implicitly chooses inventory based on the quantities it offers spot buyers.

We modify the seller's payoff in market j and month t to incorporate the net value of inventory V_{jt}:

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

The net value of inventory V_{jt} can include many factors that enter into the seller's inventory decision, such as discounted profits given market conditions and storage costs.

When the seller makes take-it-or-leave-it rebate and quantity offers to the spot buyers in $N_{jt} \subseteq S_{jy}$, it no longer faces the constraint that total quantity is at most Q_{jt}^S. However, the seller does consider the effect of an offer on the net value of inventory V_{jt}. The seller makes offers to solve the following problem:

$$\max_{\substack{\{q_{ijt}\} \\ \{r_{ijt}\}}} \sum_{i \in N_{jt}} (\overline{p}_{jt}(1 - R_{ijt}) - c_{ijt})q_{ijt} + V_{jt}\left(I_{jt-1} + Q_{jt} - \sum_{i \in N_{jt} \cup C_{jy}} q_{ijt}\right) \text{ such that } q_{ijt} \leq Y_{ijt} \text{ for all } i \in N_{jt}.$$

As before, buyers act to maximize its total payoff.

To describe the equilibrium, we assume that the net value of inventory V_{jt} is weakly concave. Concavity is a natural assumption in our context for two reasons. First, the seller stores its inventory in warehouses with finite capacity, suggesting convex storage costs. Second, the declining spot price curve suggests that the marginal discounted profit from inventoried quantity decreases. The negative correlation between inventory and minimum spot price shown in Figure OA.7 supports this assumption.

Under the assumption of weak concavity, the equilibrium is straightforward to describe. The seller orders buyers by margin $\overline{p}_{jt}(1 - R_{ijt}) - c_{ijt}$ and trades as much quantity as possible with those buyers as long as margin is greater than the marginal net value of inventory.

To estimate the net value of inventory V_{jt}, we specify a functional form:

$$V_{jt} = \alpha_j I_{jt} + (\beta_V/2)I_{jt}^2 + \epsilon_{jt}^V I_{jt},$$

where α_j is a market fixed effect and ϵ_{jt}^V is normally distributed and mean-zero conditional on market j and I_{jt}. We do not include a constant term because it is not identified.

We identify the parameters by assuming that the lowest spot price we observe in market j and month t exactly equals the marginal net value of inventory. In the model, we only know that the lowest spot price is weakly greater than the marginal net value of inventory. Thus, by assuming that the lowest spot price exactly equals the marginal net value of inventory, we overestimate the marginal net value of inventory. Insofar as inventory is a substitute to quantity contracting, this assumption is conservative for our purposes.

Under the assumption, the following equation identifies the parameters that govern V_{jt}:

$$p_{jt}^{min} = \alpha_j + \beta_V I_{jt} + \epsilon_{jt},$$

where p_{jt}^{min} is the minimum spot mill gate price in market j and month t. Because ϵ_{jt} is assumed to be conditionally mean-zero, we estimate the equation using ordinary least squares.

Table OA.6 presents the main estimates. We estimate that the marginal net value of inventory decreases at a rate of 0.06€ per ton. This coefficient is the inverse slope of the line in Figure OA.7. Note that all the estimates in the top panel are unchanged relative to the base model because the addition of inventory does not affect the estimation of those parameters. The estimate of buyer bargaining power changes slightly from 0.29 to 0.30.

Table OA.6: Parameter estimates, model with inventory

Parameter	Spot match prob. β_γ	Rebate type β_θ	Quantity type (spot) β_Y^S	Quantity type (contr.) β_Y^C	Trade type β_ϕ
Ever contract	0.63	-0.07	-0.22		-0.13
Avg. Quantity >1,000 tons	-0.02	0.50	1.11	2.05	1.18
Prioritized	0.56	-0.002	0.0004		0.07
Logistics Costs	-0.19	0.004	0.03	1.49	-1.20
Market Price	0.08	-0.056	-0.04	-0.09	0.19
Time Trend (months)	-0.01	0.03	0.006	0.006	-0.008
Market FE	Yes	Yes	Yes		Yes
Buyer-market-year FE				Yes	
Mean spot trade probability γ_{ijt}		0.66			
Buyer bargaining power δ		0.30			
Slope of marginal value of inventory β_V (€/ton)		-0.06			

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). In practice, the accounting identity $I_{jt} = I_{jt-1} + Q_{jt} - \sum_i q_{ijt}$ does not hold because of various filters we apply to the sample, such as excluding an internally-integrated buyer. We impute production Q_{jt} to satisfy the identity. In the regression that estimates β_V, we exclude two market-months where no spot buyers trade or where all spot buyers trade, since the estimating equation no longer holds in those cases. The contract selection model is not estimated as that is not necessary for the value calculation.

The addition of inventory does not change our main results, quantitatively or qualitatively. Figure OA.8 presents the histogram of the value of quantity contracts in the model with inventory. The average value is 18% of profits. This represents an increase from the baseline estimate of 10% of profits. Theory suggests that the ability to inventory should decrease the value of quantity contracts. If the seller faced no cost of inventory, then it could use inventory instead of quantity contracts to smooth variation in spot market outcomes.

The difference between the value of contracts in the base model and in the extension is small because

Figure OA.8: Quantity Contracts Are Valuable in Model with Inventory

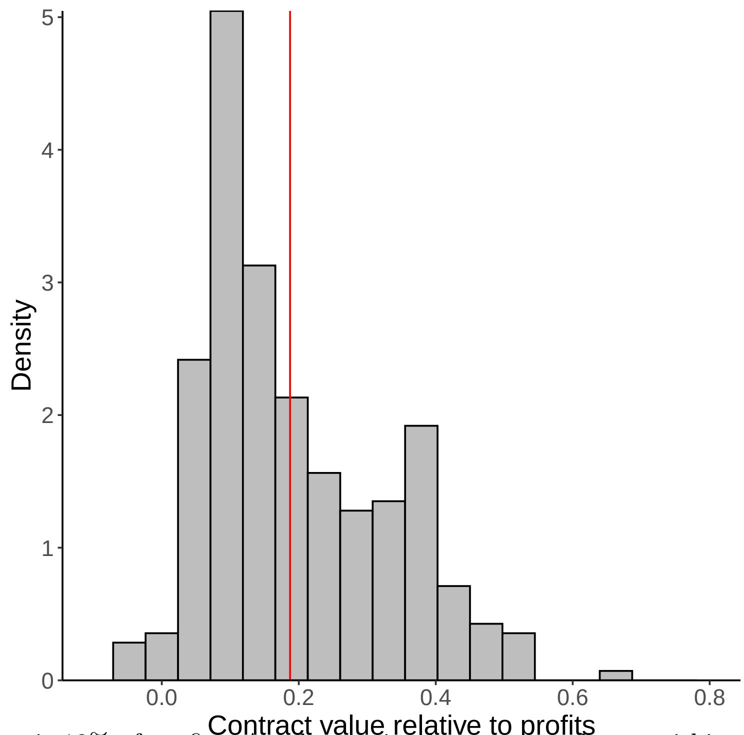

Notes. The mean value is 18% of profits. An observation is a contract buyer within a market and year.

the correlation between inventory and minimum spot price is small. This correlation informs the extent to which the seller changes inventory behavior when a contract buyer counterfactually trades on the spot. The small correlation implies that the seller does not change inventory behavior much. Thus, the addition of inventory to the model does not dramatically change our results.

Online Appendix C.2 Hold-up Analysis

We do not find empirical evidence that hold-up is an important driver of outcomes in our setting. Models featuring hold-up predict that the future value of a relationship prevents opportunism. If hold-up was a major concern in our setting, then the seller would have to compensate the buyer for an unusually high price with a high future rebate. Accordingly, we evaluate whether shocks to a buyer's contemporaneous payoff correlate with the level of the buyer's rebate the following year.

Table OA.7 presents the results of an instrumental variables regression analysis designed to quantify the correlation between a buyer's contemporaneous payoff and future rebate. The dependent variable of interest is mill gate price. We instrument for mill gate price with variable production costs taken from the invoice data. Changes in price driven by variable production cost change the buyer's payoff as long as these cost shocks are uncorrelated with other drivers of demand (conditional on controls). We take the quantity-weighted annual average of all variables. The coefficient on mill gate price in the first column is statistically indistinguishable from zero. Thus, we find no evidence of hold-up. We include the corresponding reduced form and first stage regressions as well.

Table OA.7: Small Correlation between Cost Shocks and Future Rebate

	(IV) Future Rebate	(RF) Future Rebate	(FS) Mill Gate Price
Mill Gate Price	-0.00004		
	(0.0004)		
Variable Production Costs		-0.00001	0.39
		(0.0002)	(0.10)
Market Price	0.02	0.02	35.74
	(0.02)	(0.006)	(3.19)
Logistics Costs	-0.18	-0.18	-10.24
	(0.03)	(0.03)	(13.50)
Months Since Jan. 2014	-0.00009	-0.0001	1.64
	(0.0008)	(0.0004)	(0.23)
Priority	-0.10	-0.11	75.71
	(0.03)	(0.01)	(7.68)
Observations	256	256	256
R-squared	0.35	0.33	0.71
F statistic (excluded instrument)			13.51
Sample	Contract	Contract	Contract
Market FE	NO	NO	NO

Notes. In the first column, we instrument for mill gate price with variable production costs. In the second and third columns, we present ordinary least squares estimates. Robust standard errors are in parentheses. All variables are quantity-weighted annual averages. Future rebate is defined as one minus price after rebates over market price, in the following year. Mill gate price and variable production costs are measured in euros. 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).