

Incomplete Contracts in Commodities Trade: Evidence from LNG

Swati Dhingra, Ningyuan Jia, Gianmarco Ottaviano, Thomas Sampson, and Catherine Thomas

LSE and Bocconi

CEP and CEPR

PRELIMINARY. PLEASE DO NOT CITE.

November 14, 2023*

Abstract

Long-term quantity contracts between buyers and sellers of commodities are common even in industries where goods are traded on global spot markets. This paper investigates the relative efficiency of the long-term contracts between buyers and sellers of liquefied natural gas (LNG) that accounted for 75% of total industry volumes from 2009 to 2019. The results show that the contracts are incomplete in ways that allow shipments to adapt to varying demand conditions. Estimates from a structural model of the industry—identified by exogenous weather shocks—demonstrate how adaptation within contractual relationships can be incentive compatible for both parties. Sellers receive around a third of contract surplus, which is sufficient to compensate them for forgoing spot markets and for exerting relationship-specific investment to deliver shipments that fit buyer demand. A counterfactual that supplies all shipments to spot markets, while holding quantities and final destinations constant, suggests that adaptation within contracts contributed around 30% of global industry value over the period studied.

*We thank Enrico Cavallotti and Edoardo Tolva for excellent research assistance. We are very grateful to Kpler for providing the detailed industry data. Heski Bar-Isaac, Jordi Blanes, Davin Chor, Rikard Forslid, Bob Gibbons, Erik Hurst, Rocco Macchiavello, and Saskia Meuchelböck, as well as seminar participants at Aarhus, ERWIT, LSE, Lucca IMT, MIT, and OASIS provided very helpful comments. All conclusions are our own.

1 Introduction

Long-term contracts are common in situations where buyers and sellers are locked into bilateral relationships having made transaction-specific durable investments (Klein, Crawford, and Alchian (1978), Williamson (1979), Masten and Crocker (1985)). Contracts increase relationship surplus by reducing costly future contention about how to divide the quasi-rents. However, ex ante investments with long time horizons require contractual arrangements that are flexible enough to adapt to unpredictable changes to supply and demand without being impractical to design and implement.

The paper examines the efficiency of long-term contracts in the global liquefied natural gas (LNG) industry that last up to twenty years. There are two main reasons to study this setting: First, it is a growing source of global energy, supplying around 10% of world demand, with bilateral contracts governing around three quarters of this supply.¹ Second, it provides a market where the value of customized intermediate inputs in a global supply chain can be directly estimated because weather conditions provide exogenous shocks to demand and data reveal the times that individual shipments arrive at their destinations.

Our research question is how well contracts achieve potential gains from trade by mitigating supply chain frictions, enabling supply to be timed to better match buyer demand. We also estimate the distribution of the gains between buyers and sellers, within contracts and in local spot markets.

We first describe the salient industry features that motivate our research question and guide our modeling choices. On the supply side, sellers of LNG have made large capital investments in liquefaction plants, called trains, and face very low marginal production costs. As a result, plants tend to produce at close to capacity and total output is unresponsive to spot prices. LNG is transported from sellers to buyers in large, specialized vessels that are often owned (or long-term leased) by one of the two parties. Buyers, who have also made large capacity investments in regasification terminals, and who typically operate with some excess capacity, are subject to demand shocks due to unexpected weather conditions. Contracts specify an annual quantity to be transacted between the buyer and seller at a price that is typically indexed to oil prices rather than to destination-specific LNG spot prices. Zahur (2022) analyzes the timing of contract signing relative to capital investments and documents that contracts duration can be as long as the typical productive life of the initial investments.

This paper focuses on how contracts operate from week to week, as buyers and sellers bilaterally negotiate the arrival times of individual contracted shipments throughout the year. Data on shipment times and weather conditions between 2009 and 2019 in the largest LNG destinations show that the quantity of contracted LNG arriving at a destination responds to local weather shocks. In weeks when it is unexpectedly cold, relative to seasonal norms, buyers receive more LNG under contract. That is, sellers

¹Natural gas provides around 23% of the world's energy and nearly 40% of natural gas is supplied in its liquefied form (IEA, 2021).

deliver LNG at times when it is particularly valuable to their contract partners, treating the buyers as “priority customers”. It is notable that these times coincide with when the seller’s outside option—selling on spot markets—is also more valuable because spot prices are also positively correlated with weather shocks.

Our analysis shows that delivered contract shipments are valued more than the shipments available to buyers on their local spot markets at the same time. Although LNG is itself a commodity, it is transported in large vessels that take time to travel the often long distances between sellers and buyers. The relatively small number of shipments in each spot market may not coincide with buyer demand at any point in time. We find that sellers exert relationship-specific effort within contractual transactions to deliver shipments in volumes and at times that match buyer demand. Because seller effort is costly, and sellers forgo favorable spot market outside options when they fulfil contract requests, sellers must anticipate extracting around one third of the overall estimated contract relationship value via the contract price in order for them to find contracts incentive compatible. We show that contracts contribute to total industry gains from trade. A counterfactual industry-level analysis that converts all contract shipments to spot is welfare-reducing.²

The shipment-level data used in this paper were made available to us by the company Kpler, a provider of data and analytic solutions for commodity markets. The data include information on the origin and destination of each LNG shipment, the volumes at origin and destination, the nature of the vessel, and the voyage duration. The data on LNG spot prices was downloaded from EIKON and destination-specific weather data came from the World Bank Climate Knowledge Portal.

We present three sets of empirical analyses. The first is a reduced-form investigation of variation in volumes at the route-week level. Larger quantities flow to destinations when it is unseasonably cold. We show that it is contracted volumes that respond to weather shocks and that spot volumes do not adjust. These findings are evidence that sellers try to send contract shipments to their relationship partners at times when buyers have the highest value for those shipments.

The second analysis explores the hypothesis that buyers value the contract shipments they receive over available spot shipments, given observed spot prices and quantities. The logic underlying this hypothesis is that if there were no such premium, there would be no contract relationship surplus that could be shared with the seller, and no economic rationale to explain why sellers are consistently willing to send contract shipments rather than sell on the spot market exactly when the spot prices are high. Using buyers’ revealed preference for contract shipments, we estimate a lower bound on the typical contract shipment

²Consistent with this conclusion, contractual relationships have proved to be very resilient during the energy crisis in 2022 arising because of the war in Ukraine. Sellers continued to deliver contract shipments despite all-time high spot prices. Two notable exceptions were Eni and Gunvor in early 2022, who did not deliver anticipated contract shipments to Pakistan [S&P \(2022\)](#). While the total share of spot volumes had been increasing up to 2020, this trend has recently reversed, “LNG term contract volumes have leapt this year, as energy security becomes paramount worldwide” [S&P \(2023\)](#).

value premium that is shared between the buyer and the seller. Intuitively, this premium must be sufficient to compensate sellers for not deviating from the contract when spot prices are high and compensate buyers for not deviating from the contract when spot prices are low. The estimated mean of the relative value for the contract shipments studied is \$27.07 per cubic metre of LNG, which is around 15% of the mean spot price in the sample used here.

This second analysis is limited in that it produces lower bounds of relative contract value compensating sellers only for foregoing spot markets when contract shipments are requested. It does not allow for sellers to be exerting any relationship-specific effort to increase buyer value. To ask the more fundamental question of how value is created and shared along the value chain, and how much value contracts create compared to the counterfactual of no contracts at all, we turn to the third analysis in the paper, which constructs and estimates a structural model of the LNG industry.

In the model, we allow buyer value for LNG to be a function of local weather. For contract shipments, buyer value is also a function of relationship-specific seller effort. The amount of effort the seller exerts depends on how much the effort contributes to buyer value, effort costs, and the share of contract surplus the seller extracts. In each period, buyers observe weather conditions and choose whether to request a contract shipment or to buy on the spot market. The spot value depends on an idiosyncratic shock drawn from a distribution with an expected value that is an increasing function of local spot market thickness in that period.³ The model delivers a set of theoretical moments, derived from incentive compatibility constraints, that give the probability a contract buyer requests a contract shipment in a given week as a function of the weather and model parameters.

Because the model focuses on route-week-level buyer decisions, it does not explicitly take into account seller capacity constraints or any trade-offs the seller has to make between demand from multiple contract partners in any one period. We do two things to accommodate these seller-level constraints. First, we include a fixed cost of fulfilling contract shipments that is a function of how many contract requests a seller fulfils that month relative to their typical number. Second, we estimate the model parameters on a subset of the data: route-weeks where the seller sells at least some positive quantities on the spot markets. These are quantities that sellers could have otherwise allocated to contract partners without being capacity constrained.

We estimate the model using maximum likelihood. Parameter identification is made possible by independent variation in the relationships between contract shipment timing and spot prices, holding weather constant, and between contract shipment timing and weather shocks, holding spot prices constant. Using the model estimates to predict the untargeted moments of the elasticity of contract shipments with respect

³Spot market thickness is measured as the relative size of the spot market in the region and how well the distribution of vessel sizes in the spot market fits the distribution of local contract demand.

to weather shocks and spot prices gives values that are close to those in the data.

The results include distributions of estimates of the willingness to pay for the LNG in contract and spot shipments on each route-week in the sample. A cubic metre via a contract shipment is more than twice as valuable, on average, given the current size of the spot markets.⁴ The lower bound on the share of contract shipment surplus that sellers must extract via the contract price in order for contractual relationships to be incentive compatible is estimated to be 29%. We also find that, for around one half of contract shipments, sellers exert relationship-specific effort to deliver contract shipments that buyers value more highly, for example to arrive at the ideal time during a week or in desirable quantities.

We then turn to a counterfactual analysis that asks how much long-term contracts contribute to industry value compared to a world where all LNG is traded on local spot markets. We are able to do this by holding fixed total global supply and sellers' allocations of shipments to destinations in each week. We convert all arriving contract shipments to spot and add them to the spot shipments already arriving at the destination in a given week, giving us a counterfactual spot market in each destination. Because we have estimated the relationship between local spot market thickness and expected spot value, we can simulate the expected values buyers would have obtained in these enlarged counterfactual spot markets.⁵

We find that the efficiency of the LNG industry would have been significantly reduced in the absence of contracts between 2009 and 2019. The gains from contracts outweighed the fixed and variable costs of maintaining them, contributing around 32% of the total gains from trade in the industry over this period.

In all the analyses in this paper, we take the existing contract network as fixed. There is some reduced form evidence that sellers tend to have contracts with buyers in destinations with more unpredictable weather, but we do not model why some routes have contracts and others do not. We have also not modeled transport costs explicitly.⁶ One other important issue our analysis does not consider is seller market power. In theory, it could be that sellers choose to fulfil contract shipment requests in order to withhold quantities from spot markets and gain from the resulting higher spot prices. We do not allow sellers' supply decisions to include these strategic considerations. The main reason for this choice is that the routes between sellers and buyers tend to either be contract or spot. That is, if a seller has a contract with a buyer, it is unlikely that they are also selling on the spot market at that destination, and vice

⁴We note that a part of this difference arises because buyers tends to request contract shipments during weeks when they experience positive demand shocks.

⁵This approach allows us to sidestep the challenge of estimating counterfactual spot prices because we focus on the total surplus created by each shipment as a function of weather conditions, rather than how the unobserved counterfactual spot price divides this surplus. Setting up the counterfactual in this way does not take into account any allocative efficiency gains possible in the spot market if sellers were to divert spot shipments to locations where they anticipate higher prices.

⁶The counterfactual estimation holds the distances traveled by each vessel fixed and, as such, transport costs play no role in the welfare calculations.

versa.⁷

This paper relates to several literatures. First, are industry-specific studies of the pros and cons of long-term quantity contracts. These include [Joskow \(1988\)](#), and [Pirrong \(1993\)](#), and the discussion in [Lafontaine and Slade \(2012\)](#). [Zahur \(2022\)](#) has studied the LNG industry in depth and has focused on the long-term contracts that we investigate here. He shows that contracts allow greater capital investments at the cost of allocative efficiency but does not examine whether contractual relationships create value through week-to-week adaptation to shocks. Our analysis is motivated by Zahur’s observation that contracts are often signed around the time of the capital investments, which, due to the long life of capital, raises the issue we explore here of the efficiency implications of contract flexibility in delivery timing.⁸

In our setting, unpredictable weather causes exogenous demand shocks that we exploit to identify how supply chains respond. We establish that long-term contracts are sufficiently flexible to allow the supply side to adapt and add to overall industry value. This finding complements work showing that vertical integration can permit adaptation to unforeseen circumstances. [Forbes and Lederman \(2009\)](#) and [Forbes and Lederman \(2010\)](#) study vertical integration in the airline industry. They show that major airlines tend to own regional airline partners when adverse weather is more likely and where a route is more integrated into the major’s network. Integrated routes also respond to weather shocks and perform better than non-integrated routes at the same airport on the same day. [Costinot, Oldenski, and Rauch \(2011\)](#) show that the boundaries of U.S. multinational firms reflect the extent to which sectors involve non-routine tasks. Vertical integration along the supply chain is more likely in sectors with high problem-solving intensity, consistent with ownership facilitating adaptation to changing circumstances. We find that long-term contracts enable adaptation in LNG supply chains in terms of delivery volumes and timing. We highlight that this finding implies contracts increase the resilience of the industry supply chain to demand shocks.⁹

A recent literature on relational contracts establishes how the prospect of future rents effectively deters short term opportunism in a range of empirical settings ([Macchiavello \(2022\)](#)). [Macchiavello and Morjaria \(2015\)](#) use the dynamic incentive compatibility constraints implied by repeated seller-buyer interactions to estimate lower bounds on the value of relationships in the Kenyan flower industry, a setting of imperfect contract enforcement. We show that the future value of formal, but incomplete, contracts in the global LNG industry governs the short term actions of the sellers in the industry. Buyers, however, compare contemporaneous costs and benefits, interpreting contracts as call options to request shipments at times

⁷While there are several large firms in this industry that have minority ownership stakes in multiple sellers, the majority ownership stake is often a state energy company with no cross-seller links.

⁸Because our counterfactual does not vary each shipment’s destination, we do not explore allocative efficiency.

⁹Recent global events, perhaps most notably the Covid-19 pandemic, have prompted much political debate and study of the efficiency implications of potential supply-side disruptions within global value chains (see [Grossman, Helpman, and Lhuillier \(2023\)](#), [Grossman, Helpman, and Sabal \(2023\)](#), and the papers cited therein).

when they are particularly valuable.

One recent industry-specific study that examines the externalities that contracts exert on spot markets is [Harris and Nguyen \(2022\)](#).¹⁰ They estimate the incentives that govern both relationship formation and the actions of buyers and sellers in U.S. trucking. Relationships form between trading partners with a high match quality, but the large number of contract relationships leads to thin spot markets, increasing the search costs for trading partners. [Tolvanen, Darmouni, and Essig Aberg \(2021\)](#) also find that frictions in spot markets provide an insurance role for buyers, in analysis of transactions data from a large producer in the U.S. pulp industry. In our setting, we show that relationship surplus is an endogenous outcome reflecting demand shocks and seller incentives to exert costly relationship-specific effort to increase buyer value. Even though these costs are saved in the counterfactual that eliminates contracts, the increased thickness of the spot market is insufficient to compensate for the lost returns of seller effort.

Our interpretation of the relatively low buyer values for spot market shipments is that spot markets do not supply the quantities of LNG that buyers need at the ideal delivery times, even in the counterfactual thicker spot markets, due to frictions related to bulky shipping technology and the long distances between buyers and sellers. [Brancaccio, Kalouptsi, and Papageorgiou \(2020\)](#) document how frictions arise in global shipping when country-level trade imbalances in dry bulk commodities lead to shipping price differences across markets. We note that LNG vessels are specialized and there is no two-way trade. In addition, vessels are very large relative to total weekly demand, an indivisibility that amplifies the distance-related time lag between requesting and receiving a shipment. We infer that contracts were efficient in this industry because the total market size was too small, over the time period studied, for spot markets to overcome these geographical and technological frictions.

This paper proceeds as follows: Section 2 describes the industry and summarizes the data. Section 3 is the reduced-form analysis relating contract volumes to weather conditions. Section 4 estimates lower bounds on buyers' values for contract shipments given prevailing spot prices. Section 5 presents the structural model, Section 6 discusses estimation, and Section 7 presents the results. Section 8 is the counterfactual analysis. Section 9 concludes.

2 Industry and Data

2.1 Industry overview

The global LNG industry allows the gas produced in dispersed, oil-rich, locations to be liquefied by cooling it to -160° centigrade and transported via ships to locations with regasification capacity. The industry

¹⁰[Kranton \(1996\)](#) sets out theoretically how spot market externalities can be large enough or small enough to determine whether relationships are efficient.

is characterised by large, long-lived, location-specific capital investments at both origins and destinations typically located far apart from each other, and connected via a costly shipping technology. Demand is concentrated in East Asian countries, although regas terminals continue to be built at ports worldwide.¹¹

Long-term sales and purchase agreements (SPAs) are contracts between sellers and buyers that govern the bulk of delivered volumes and that are signed before or after capital investments are made (Zahur (2022)). Contracts specify annual volumes and a pricing formula. Much of the detail in the contracts relates to their “take or pay” nature, a feature shared with pipeline gas contracts (Masten and Crocker (1985)). Prices are often indexed to local oil prices, which means that the seller bears risk related to the forgone opportunity of selling at high LNG spot prices when they occur. The buyer typically bears the quantity risk, in that there is an obligation to take the total quantity over the course of the year, whether they need it or not, or incur demurrage charges to compensate the seller. Contracts also contain various clauses related to unforeseen contingencies and the length of acceptable delivery windows.

The SPAs, however, do not specify the exact timing of individual contract shipments during the year. While parties typically operate under a three-month forward-looking schedule, shipments are governed by separate Confirmation Notices that describe timing, quantity, and other shipment-specific details. Confirmation Notices are agreed bilaterally prior to shipping. Therefore, the governing contracts are incomplete in that sellers and buyers must arrange exactly when and in what quantities each shipment is delivered (or collected, when the buyer is responsible for the shipping). Because buyers are subject to varying weather conditions that impact shipment value, and shipment delivery dates are observable, this industry offers a setting where we can observe if contractual partners are able to arrange transactions at valuable times.¹²


Residual LNG volumes are traded on global spot markets. There is significant variation in the destination-specific local spot price over time, and also across countries at any one point in time. Figure 1 shows the price data we use in this paper.

Appendix A shows that the largest exporters in 2019, the last year of data used in this study, were Qatar, Australia, the USA, and Russia. Malaysia, Nigeria, and Indonesia, the three next largest suppliers, jointly supplied less than each of Qatar and Australia. Australia’s destinations were concentrated in Japan and China, whereas Qatar’s were more spread out. LNG is relatively inelastically supplied up to capacity

¹¹Recent disruptions to pipeline gas in the EU due to the Russia-Ukraine war have led to many new regasification investments. The import capacity in the EU and the UK is set to increase by one third by 2024, after expanding only modestly in the 10 years to 2022. <https://www.eia.gov/todayinenergy/detail.php?id=54780>.

¹²Although contractual agreements are proprietary, template SPAs are publicly available. For example, the company BP publishes its master sales and purchase agreement template on its website, [https://www.bp.com/content/dam/bp/business-sites/en/global/bp-trading-and-shipping/documents/bp-master-ex-ship-lng-sale-and-purchase-agreement-2019-edition%20\(3\)%202.pdf](https://www.bp.com/content/dam/bp/business-sites/en/global/bp-trading-and-shipping/documents/bp-master-ex-ship-lng-sale-and-purchase-agreement-2019-edition%20(3)%202.pdf). A Confirmation Notice template for a specific shipment is given in the Appendix, Schedule 1, on pages 46 to 49.

Figure 1: LNG spot prices by destination country over time, USD



at each liquefaction plant once capital investments have been made. Every now and again there are planned or unplanned supply outages due to routine maintenance or unplanned events. There is much less global liquefaction capacity than global regasification capacity.

2.2 Data

We have shipment-level data on 44,848 shipments from January 2009 to December 2019 that were either contract or spot market transactions. 13 countries imported more than 500 shipments in total over the 11 years in the data. Japan has been, by far, the largest importer historically, receiving over 15,000 shipments in total during the 11 years. South Korea had the second highest number of shipments, at nearly 6,000. China has steadily increased shipment imports, from around 100 in 2009 to nearly 1,000 in 2019. For each shipment, our data contain distance traveled, time taken, vessel capacity and type, as well as volume at origin and volume at destination, which permits a boil off calculation.¹³

We have 36 origins in the data that we currently use. Of these, 17 are countries where liquefaction is under control of one (often state-owned) company, and integrated oil companies often have minority stakes in these state-owned oil companies. The countries are: Algeria, Angola, Brunei, Cameroon, Egypt, Equa-

¹³The volume arriving at the destination is always lower than the volume that left the origin. A note on conversion factors. To convert mmBtu into cubic metres of LNG: $1mmBtu = 0.048$ cubic metres of LNG. Equivalently, $1m^3 = 21.04mmBtu$. Source: <http://www.lngplants.com/conversiontables.html>

torial Guinea, Indonesia, Malaysia, Nigeria, Norway, Oman, Papua New Guinea, Peru, Qatar, Trinidad and Tobago, United Arab Emirates, and Yemen. The other 19 origins are separate liquefaction facilities in Australia (10), Russia (3), and the USA (6). We treat these plants separately because they have different owners and/or geographies: The Russian plant Sakhalin is in the far East, very near Japan, and Yamal and Vysotsk are in the West of Russia. The Australian plants are also located on different sides of the country, as well as being separate companies, albeit with some overlapping ownership.

We have 39 destinations, which are the regasification terminals aggregated to the country level. The largest individual buyers are often state or city energy providers. For example, the largest single buyer in the data is the only South Korean buyer, KOGAS, a public natural gas company. The next largest is CPC Corporation, the sole LNG importer in Taiwan, which is a state-owned company. CNOOC, the state owned China National Offshore Oil Corporation, is the largest importer in China. In Japan, the three largest importers are the Tokyo Electric Power Company (TEPCO), which is majority-owned by the government of Japan, and its joint venture JERA, as well as the Tokyo Gas Company.

We hence have 1,404 possible origin-destination pair routes in each of 11 years, but, since some facilities came online during the sample, we have 8,570 route-years pairs where the plant and terminal are both active during the year. Of all feasible route-years, only 2,178 contain any positive flows.

In the empirical analysis, we restrict attention to the 10 destinations with large LNG spot markets, where weekly spot price data is available. These are the destinations shown in Figure 1.¹⁴ These 10 destinations account for 31,267, or 70% of all the contract and spot shipments over the time period. The shipment-level data allow us to see the total volume at origin and at destination traded on a given route in a given week by summing up the shipment volumes. Because we can see whether any given shipment is contract or spot, we can also compute the total contract and spot volumes.

We use data from several other data sources. The location-specific LNG spot prices shown in Figure 1 are from EIKON. These are weekly data for the 10 different spot markets we study. Country-specific temperatures are from the World Bank climate knowledge portal.¹⁵ We select the mean monthly temperature in each country for each year.¹⁶

In this paper, we are interested in how well contracts allow parties to adapt to unforeseen circumstances, in our case, unexpected weather shocks. We follow the climatology literature (see, for example, [Nicholson](#)

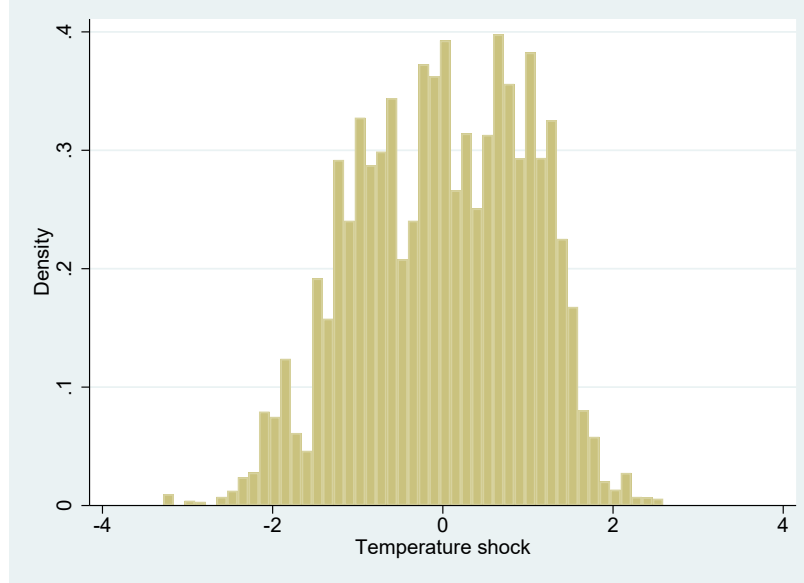
¹⁴The countries are Belgium, China, France, India, Japan, South Korea, Spain, Taiwan, Turkey, and the United Kingdom. There is a spot price also in the US, but we have excluded the US as a destination as it is predominantly an exporter of LNG. We also exclude the LNG spot markets in Argentina, Brazil, and Mexico. This is because they are relatively small markets and also because demand appears not to vary with weather. Also note that we assume buyers in Taiwan can purchase on the China spot market due to its proximity.

¹⁵<https://climateknowledgeportal.worldbank.org/download-data>

¹⁶All the results are robust to using the minimum monthly temperature in each year to measure weather shocks.

(1993)) and compute weather anomalies as the deviations from the destination’s long-term mean, divided by its long-run standard deviation. To adjust for seasonality, we find this measure of the shock for each destination-month. That is, the temperature shock is the actual mean temperature in a destination that month minus the mean mean temperature for that destination-month divided by its standard deviation. Because each observation is standardized, the average value of this variable is zero and the standard deviation is one. The variation in weather anomalies is shown in Figure 2.

Figure 2: Destination-month temperature shocks, standardized



2.3 Descriptive statistics

We first look at which 2,178 of the 8,570 potentially active route-years actually have some positive flows. Table 1 examines whether flows follow gravity. It contains three regressions of the indicator for a route-year being active on the log of the direct distance (in '000km) between the origin and destination and variables relating to the total volume produced (received) by the origin (destination). Column 1 shows that an increase in route distance of 1,000 km lowers the probability of the route-year being active by 20% (the coefficient is 0.204). Column 2 includes the total volume of the origin and destination as controls, and column 3 includes origin and destination fixed effects. In this final column, an increase in distance of 1000 km reduces the probability of being active by 16%. In other words, distance matters in that having any positive flows follows gravity.

Turning to the set of route-years with positive flows, the volumes on a given route in the year tend to be either contract or spot volumes. Figure 3 shows the bi-modal share of volumes on a route-year that is under contract. Based on this figure, we construct an indicator variable for whether an active route-year is mostly under contract, defined as having at least 50% of total flows in the year being contract flows and investigate selection into being a contract route.


Table 1: Probability route active in a year.

	(1)	(2)	(3)
lnairlinedistancekm	-0.204*** (0.007)	-0.187*** (0.006)	-0.164*** (0.006)
totalvoloryear		0.003*** (0.000)	
totalvoldestyear		0.005*** (0.000)	
Observations	8,570	8,570	8,570
R^2	0.103	0.305	0.525

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Figure 3: Share of route-year volumes that is under contract



In Table 2, the binary dependent variable is equal to one if the volumes on a route-year are more than 50% contract and all columns include year fixed effects. Column 1 shows there is no evidence that distance matters, in that shorter routes are no more or less likely to be contract routes in any year. Instead, the total volumes of the buyer and seller are positively associated with being a contract route. Column 2 shows that route-years originating at larger sellers and ending at larger buyers are significantly more likely to be contract route-years. Column 3 includes buyer fixed effects and shows that the typical buyer is more likely to have contract routes with its larger sellers. Column 4 includes seller fixed effects and shows that the typical seller is more likely to have contracts with its larger buyers.

Column 5 of Table 2 presents some evidence that routes select into having contracts when the weather at the destination is more unpredictable. It controls for year and seller fixed effects and the dependent variable is the mean absolute value of the monthly weather abnormality at the destination during that year, where weather abnormality is the standardized temperature shock shown in Figure 2. The mean of this measure is 0.663 and the standard deviation is 0.598. The coefficient of 0.299 suggests that a one standard deviation decrease in route-year weather predictability (an increase in unpredictability) is associated with an 18% increase in the probability that a given route-year is mostly contract. While our structural model will take the existence of a contract on a route as given, the evidence in column 5 is consistent with there being a contract in place when the potential gains from adaptation to weather shocks are largest.

The empirical analysis in the paper focuses the volumes and number of shipments at the route level in a given week. In the 2,178 route-years with some positive flows, there are 110,745 route-week observations. We match 52,201 of these route-week observations to a destination-level LNG spot price because they are in one of the 10 locations we consider. Of these, only 17,516 route-weeks have positive volumes. That is, the route-weeks with zero shipments are important in the data.

Table 3 summarizes the data used in Section 3, which are volumes of LNG at the route-week level between 2009 and 2019. The first row of column 2 shows that the mean volume per route-week is 82.75 thousand cubic meters. Conditional on there being some positive volumes, the mean volume is 246.61 thousand cubic meters, and the mean number of shipments on a route-week is 1.79. The mean vessel capacity in the shipment-level data is 130 thousand cubic meters.

The second row of Table 3 shows that 65% of route-week observations are on routes where at least 50% of volumes that year are under contract. Hence, 35% of routes are not contract routes, and the majority of volumes on these routes is spot.

The final three rows of Table 3 show that the mean spot price for the route-weeks in the data is \$182.40 per cubic metre, which corresponds to \$8.67 per MMBtu. The mean monthly temperature across the 10 destinations is 12.77 degrees centigrade, and the standard deviation is large. The final row summarizes the weather shock used in the data (see also Figure 2).

Table 2: Probability active route-year is contract.

	(1)	(2)	(3)	(4)	(5)
Distance, ln(km)	0.017 (0.015)	0.006 (0.015)	0.005 (0.016)	-0.014 (0.015)	
Seller volume, year		0.001*** (0.000)	0.001** (0.000)		
Buyer volume, year		0.001*** (0.000)		0.002*** (0.000)	
Buyer-year weather shock					0.299*** (0.044)
Constant	0.566*** (0.027)	0.528*** (0.029)	0.566*** (0.031)	0.553*** (0.026)	0.432*** (0.031)
Observations	2,171	2,171	2,169	2,171	1,449
R^2	0.011	0.021	0.130	0.273	0.289
Fixed effects	Year	Year	Year, buyer	Year, seller	Year, seller

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 3: Active route-weeks, 2009 to 2019

	Volumes, '000 cubic metres				
	Obs	Mean	Std.dev	Min	Max
All	52,201	82.75	152.25	0	1904.41
Contract	34,144	110.98	173.81	0	1904.41
Non-contract	18,057	29.37	74.95	0	722.51
Local spot price, USD per m^3	52,201	182.40	82.19	47.34	426.06
Local mean temp, °C	52,201	12.77	9.24	-8.99	31.25
Local temp shock, °C	52,201	0.00	1.00	-3.29	2.59

3 Reduced form analysis

This section asks if route volumes respond to weather shocks at the destination, and if the responsiveness varies for contract and spot volumes. We consider only active routes, defined as those where plants at both the origin and destination are operational and there is at least one shipment on that origin-destination route during the calendar year. The weather shock variable is as defined in Section 2.3, as the mean monthly temperature at the destination standardized by subtracting the mean of that destination-month’s mean temperature and dividing by its standard error all years in the data, and shown in Figure 2.

In the first set of specifications, we regress total volumes on a given route-week, in ‘000 cubic meters, on the weather shock that month, as follows:

$$Q_{lgt} = \alpha w_{gt} * O_{lg} + \Gamma_{gm} + \epsilon_{lgt}, \quad (1)$$

where Q_{lgt} is the total volume on active route lg in week t , w_{gt} the weather shock that month, and Γ_{gm} are destination-month fixed effects. The variable O_{lg} , included only in the second specification, is an indicator variables for how the route is organized, equal to 1 for contract routes.

The results of estimating equation (1) are given in Table 4. Column 1 shows that the total volume is higher by 2.820 thousand cubic meters when the weather shock is lower by one standard deviation. This corresponds to a 3.4% increase in total volume relative to the mean. Column 2 includes the interaction of temperature with the indicator for whether the route is a contract route. The coefficient on non-contract routes becomes a small, positive and significant number, that is, volumes actually decrease on non-contract routes when there is an adverse weather shock (a positive demand shock). For contract routes, the sum of the relevant coefficients gives that a one standard deviation negative weather shock is associated with a 4.11 thousand cubic metre increase in volumes on the typical contract route. This is a 3.7% increase in total volume relative to mean contract route volumes.

The second set of reduced form specifications changes the dependent variable to be total volumes, contract volumes, or spot volumes on a route-week, controlling for the average quantity with route-month fixed effects. The estimated equation is:

$$Q_{lgt} = f(w_{gt})' \alpha + \Gamma_{lgm} + \epsilon_{lgt}, \quad (2)$$

where Q_{lgt} are the volumes in question, and Γ_{lgm} are the route-month fixed effects, rather than the destination-month fixed effects included in Table 4. The results are given in Table 5. In columns 1, 3, and 5, $f(w_{gt}) = w_{gt}$ and α is a scalar, investigating whether there is a linear relationship between the weather shock and volumes. In columns 2, 4, and 6, $f(w_{gt}) = w_{gt} + w_{gt}^2$ and α is a vector of two coefficients, allowing for a quadratic relationship.

Because there are route-month fixed effects in these specifications, the coefficients reflect variation in volumes delivered in the same month across years on a given route. In columns 1 and 2, the dependent

Table 4: Route-week volume responsiveness to local temperature shock.

	(1)	(2)
	volordestweek	volordestweek
Weather shock	-2.820*** (0.725)	1.555* (0.805)
Contract route * Weather shock		-5.666*** (1.338)
Contract route		78.826*** (4.648)
Constant	82.750*** (0.000)	31.144*** (3.044)
Observations	52,201	52,201
R^2	0.094	0.152
Fixed effects	g-month	g-month

Standard errors in parentheses, clustered at buyer-month level.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 5: Route-week volumes responsiveness to local temperature shock.

	(1)	(2)	(3)	(4)	(5)	(6)
	Total	Total	Contract	Contract	Spot	Spot
Weather shock	-1.319** (0.565)	-1.518*** (0.569)	-1.522*** (0.505)	-1.648*** (0.509)	-0.093 (0.237)	-0.116 (0.243)
Weather shock squared		-0.892* (0.496)		-0.566 (0.426)		-0.101 (0.214)
Constant	82.766*** (0.000)	83.656*** (0.495)	65.849*** (0.000)	66.413*** (0.425)	11.997*** (0.000)	12.098*** (0.213)
Observations	52,187	52,187	52,187	52,187	52,187	52,187
R^2	0.624	0.624	0.643	0.643	0.177	0.177
Fixed effects	lg-month	lg-month	lg-month	lg-month	lg-month	lg-month

Standard errors in parentheses, clustered at route-month level.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

variable is the total volumes on a route-week. The negative significant coefficient of -1.319 tells us that volumes are around 1.6% lower when the temperature is one degree higher. Column 2 shows that the relationship is non-linear, in that responsiveness is increasing in the weather shock. Columns 3 and 4 have the contract volumes as the dependent variable, and the coefficients suggest a negative linear relationship between the weather shock and volumes. Columns 5 and 6 find no significant relationship between spot volumes and the weather shock.

As some additional evidence that weather shocks are indeed shocks to local demand at the destination, we show that spot prices vary with the temperature shock. We regress the weekly local spot price in a destination-week on the weather shock. We have 4,423 destination-week observations and include destination-month fixed effects to control for predictable seasonality at each destination. The results are in Table 6. The estimates suggest a negative linear relationship—spot prices are high when the destination is unexpectedly cold.

Table 6: Destination spot price responsiveness to local temperature shock.

	(1)	(2)
	Local spot price	Local spot price
Weather shock	-12.419*** (2.175)	-12.367*** (2.244)
Weather shock squared		0.176 (1.388)
Constant	180.682*** (0.124)	180.499*** (1.441)
Observations	4423	4423
R^2	0.161	0.161
Fixed effects	g-month	g-month

Standard errors in parentheses, clustered at g-month level.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

4 Lower bounds on buyer values for contract shipments

This second empirical analysis estimates lower bounds on the value of receiving LNG in a contract shipment rather than a spot shipment, given the thickness of the spot market and the status quo spot prices. Our approach builds on the intuition developed in Section 3 that the buyer requests contract shipments in

periods when they have positive demand, which is also when spot prices tend to be high. We add to this the constraint that a contract relationship is also incentive compatible for the seller over the duration of a contracting period, here assumed to be a calendar year.

We first introduce some notation, which will also be used in the model in Section 5. When there is a contract in place on a given seller-buyer route, the downstream buyer, g , has a value V_{lgt} for one cubic metre of LNG arriving from seller l in a contract shipment at time t . The value g has for the same LNG quantity in a spot shipment in t is V_{gt}^s , which is not l -specific.¹⁷ The per cubic metre prices for LNG arriving via contract and spot are denoted p_{lgt}^c and p_{gt}^s , respectively. A seller can deliver the shipments they have at time t to any of their contract partners or sell on any spot market, so the relevant spot price for the seller is the highest available spot price in t . We denote this p_{lt}^s . The object of interest in this section is $V_{lgt}^c - V_{gt}^s$ and the challenge is that only p_{gt}^s are observed in the data.

We start by considering the buyer's surplus from a contract shipment, $V_{lgt}^c - p_{lgt}^c$, and appeal to revealed preference: whenever the buyer receives a contract shipment, it must create more value for them than would a spot shipment at the local prevailing spot price. That is, buyer incentive compatibility requires:

$$V_{lgt}^c - p_{lgt}^c \geq V_{gt}^s - p_{gt}^s,$$

or, equivalently,

$$V_{lgt}^c - V_{gt}^s \geq p_{lgt}^c - p_{gt}^s. \quad (3)$$

Turning next to the seller's incentive compatibility constraint, we argue that sellers must, on average, receive at least as much as what they could earn selling the same quantities on spot markets over the contract relationship. For simplicity, in this section, we let the seller be forward-looking only over the month of a relevant contract shipment. The seller's dynamic incentive compatibility constraint implies that the contract prices they receive in expectation:

$$E(p_{lgt}^c) \geq E(\max_G(p_{gt}^s)), \quad (4)$$

where expectations are taken over all spot markets, G , over the contracted time period. We note that this is certainly a lower bound on the surplus required to align sellers' incentives in the relationship since it does not account for any variable or differential fixed costs incurred on contract shipments.

If both parties' expectations are consistent, we can combine the incentive compatibility constraints (3) and (4) to give:

$$E(V_{lgt}^c - V_{gt}^s) \geq E(p_{lgt}^c) - p_{gt}^s \geq E(\max(p_{gt}^s)) - p_{gt}^s. \quad (5)$$

The right hand side of inequality (5) can be calculated using data on spot prices at all destination at times when contracts are received by buyer g .

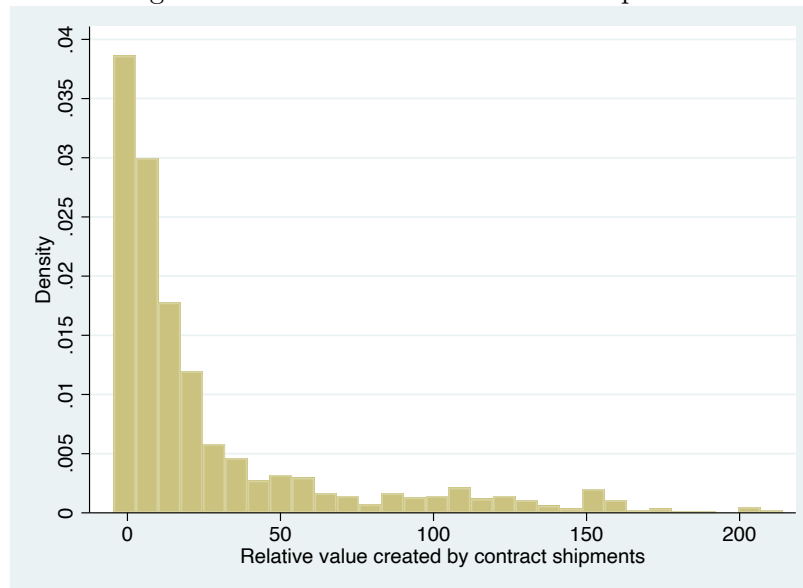
¹⁷Note that s is used to denote a spot shipment and not a seller.

The data yield an estimate of a lower bound on $E(V_{lgt}^c - V_{gt}^s)$ per cubic metre of LNG received under contract by destination g in month t , which is the contract period under consideration in this section. We find estimates for all g -months where buyer g receives at least one contract shipment.

While we have made some strong assumptions to derive these lower bounds on the value contract shipments create for buyers under prevailing spot prices, the logic follows the simple intuition that the relative expected value from a contract is sufficient to compensate the seller for not switching to the spot market when the spot price is high together with the amount required to compensate the buyer for not switching to the spot market when the price is low at times of similar demand.

Figure 6 is a histogram of the estimated lower bounds of the relative value per cubic metre of LNG arriving at the destination in contract shipments compared to available spot shipments for the 1,158 buyer-months in this analysis. The mean value is \$27.07, the standard deviation is \$40.77 and the median is

Figure 4: Relative value of contract shipments



\$10.02. A regression of destination-month value on 11 year fixed effects has an adjusted R-squared of 0.24. Including only the 10 buyer fixed effects gives an adjusted R-squared of 0.36. A regression with fixed effects for the interaction of buyer and year fixed effects gives an adjusted R-squared of 0.83. This high value comes about because the approach suggests the contract shipments dispatched when there are high spot prices somewhere in the world must be particularly valuable.¹⁸

The estimated lower bounds can be interpreted as the value of a contract shipment over a unilateral deviation to a spot shipment in the same week, conditional on the vector of global spot prices at that time. In order to derive estimates of the relative value, rather than lower bounds, we need a structural

¹⁸A similar exercise grouping time periods by destination-specific demand shock and assuming the difference between contract shipment value and contract price is at least as large as the difference between the highest and lowest prices paid for spot shipments under similar demand conditions yields similar estimates of relative contract value lower bounds.

model that delivers estimates of LNG demand and the relationship between spot market size and buyer spot value, as well as the nature of seller effort in contractual relationships and the division of contract surplus. This is the purpose of the Section 5 where we exploit exogenous weather shocks as well as local spot price variation.

5 Model

5.1 Set up

Time is discrete and each period t denotes a week. There are two types of agents, L sellers, $l \in 1, 2, \dots, L$, and G buyers, $g \in 1, 2, \dots, G$. All agents are risk neutral and have a discount factor δ .¹⁹ The buyers and sellers are globally dispersed, and vessels containing LNG travel from sellers to buyers in each period. All possible lg pairs are feasible trading routes. We assume shipping capacity is not a constraint.

Each seller l has a fixed installed liquefaction capacity and produces an exogenous quantity of LNG in each t . Each buyer g has fixed installed regasification capacity. Buyer demand for LNG varies across time, and per period demand is stochastic.

In $t = 0$, l and g decide whether to write a long-term contract specifying an annual quantity of LNG to be shipped from l to g and a formula that determines the contract price P_{lgt}^c . We assume whether or not there is a contract on a given route is exogenous to factors in the model.^{20,21} Hence, at the start of each time period $t > 0$, each l has a pre-determined set of contract buyers and each g has a predetermined set of contract sellers. Each contract is incomplete in that the exact timing and quantity of any one delivery are not contracted and shipments are agreed bilaterally throughout each year of the contract.

Transactions on any lgt route-week can take one of two organizational forms, via a contract shipment or via a spot shipment. For each period $t > 0$, the timing is as follows:

- Each buyer g :
 - experiences a demand shock and learns the nature of its demand for LNG in that period.²² All parties form expectations of the vector of all gt -specific spot prices p_{gt}^s based on observed shocks across all g .

¹⁹We do not explicitly model the shipping stage of the value chain in the current version. Specialized LNG vessels are usually owned by or on long-leases to sellers or buyers.

²⁰Selection into having a contract on a route appears consistent with reduced form evidence about the adaptation benefits of contracts. As shown in Table 2, a seller is more likely to have a contract route with a buyer with less predictable weather.

²¹Each contract's price is indexed to the destination region's oil price, P_{gt}^o , but the exact indexing rule is unobserved. Actual contract prices are also unobserved in our data.

²²We have in mind that the buyer learns the ideal day of delivery and ideal volume it needs, if any, in week t .

- makes a draw from a distribution that determines how well a shipment bought in the local spot market in time t will fit its demand conditions. The shape of the distribution reflects the expected t -specific thickness in g 's local spot market.
 - decides whether to request a contract shipment from each of its contract partners.
- Then, each seller l :
 - chooses whether to fulfil each contract request received, and, conditional on fulfilling a request, decides how much costly relationship-specific effort to exert to ensure vessel capacity and arrival time meets buyer needs.
 - dispatches contract shipments and sends any remaining output to the spot market with the highest expected price.
 - Destination spot prices p_{gt}^s are realized, equilibrating the residual spot demand and supply in each destination, and payments for contract and spot shipments are made.

5.2 Contract shipments

When l fulfils g 's request for a contract shipment in t , g 's value for the shipment, V_{gt}^c , depends on their local demand shock, w_{gt} , and on the amount of relationship-specific effort made by the seller to fit the shipment to g 's needs, $h_{lgt} \geq 0$. We specify a functional form for V_{gt}^c :

$$V_{gt}^c(w_{gt}, h_{lgt}) = v(w_{gt})(1 + h_{lgt})^\mu \quad (6)$$

In this “contract value production function”, $\mu \leq 1$ is the intensity with which l 's effort creates value. This parameter also captures any returns to scale in l 's effort.

There is a constant marginal cost of seller effort, c . Each contract shipment also incurs a fixed cost F_{lt} . We allow F_{lt} to be an increasing function of the number of contract shipments fulfilled by l in t relative to l 's typical number of contracts. This flexibility is intended to capture any increasing cost of contract fulfilment that is not g -specific. For example, perhaps it is harder to find the ideal vessel and send it at the ideal time for each contract shipment when l has more to manage.

An endogenous fraction β of contract value, net of fixed cost, goes to l and is built into the pre-agreed contract price index formula. Hence, for every contract shipment, the first order condition determining l 's level of relationship-specific effort is:

$$\beta v(w_{gt})\mu(1 + h_{lgt})^{\mu-1} = c \quad (7)$$

Rearranging equation (7) and simplifying gives the seller's choice of effort (plus one) as:

$$(1 + h_{lgt}) = v(w_{gt})^{\frac{1}{1-\mu}} \left(\frac{\beta\mu}{c} \right)^{\frac{1}{1-\mu}}. \quad (8)$$

Seller effort is increasing in effort intensity, μ , and decreasing in marginal effort costs, c . A larger share of surplus going to l , β , also leads to higher levels of effort. Note that the production function implies potential complementarities between the responsiveness of buyer value to the gt demand shock, l 's effort choice, and l 's share of surplus.

This gives variable effort costs, borne by l , as:

$$c(1 + h_{lgt}) = \beta\mu v(w_{gt}) \left(v(w_{gt})^{\frac{\mu}{1-\mu}} \left(\frac{\beta\mu}{c} \right)^{\frac{\mu}{1-\mu}} \right). \quad (9)$$

The contract shipment surplus, defined as the value to the buyer less the fixed cost, in terms of the model parameters, the endogenous variable β , and the function $v(w_{gt})$, is:

$$S_{lgt}^c = v(w_{gt}) \left(v(w_{gt})^{\frac{\mu}{1-\mu}} \left(\frac{\beta\mu}{c} \right)^{\frac{\mu}{1-\mu}} \right) - F_{lt} = v(w_{gt})^{\frac{1}{1-\mu}} \left(\frac{\beta\mu}{c} \right)^{\frac{\mu}{1-\mu}} - F_{lt} \quad (10)$$

Note that S_{lgt}^c is deterministic conditional on demand shocks w_{gt} and model parameters. It is increasing in g 's demand shock.²³ Contract shipment surplus is also increasing in l 's effort intensity and in the share going to the seller, and decreasing in the marginal cost of seller effort and in fixed cost.

5.3 Spot shipments

When g buys from the spot market in t , its value for the spot shipment is V_{gt}^s , and it has to pay the local spot market price, p_{gt}^s . We assume there is no relationship-specific investment made by the seller for these shipments. For simplicity, we also assume spot shipments do not incur any fixed cost.

Shipments in the spot market are differentiated by vessel type and arrival time, which affect spot market value. The gt -specific shock to the spot value is ϕ_{gt}^s :

$$V_{gt}^s = \frac{v(w_{gt})}{\phi_{gt}^s}. \quad (11)$$

We let $\phi_{gt}^s \geq 1$ follow a Pareto distribution \mathcal{F}_ϕ with pdf $f_\phi = \frac{\lambda_{gt}}{\phi^{\lambda_{gt}+1}}$. We allow the λ_{gt} parameter to be an increasing function of the thickness of the local spot market in gt . This captures the idea that g is likely to find a better fit for its needs when there are relatively more spot shipments available and when those shipments' capacities are similar to the capacities of the contract shipments that are typically received.²⁴ The expected value of $\left(\frac{\lambda_{gt}}{\phi^{\lambda_{gt}+1}} \right)$ is an increasing function of λ_{gt} . Larger realized draws of ϕ_{gt}^s mean that g can find only a poor match for its needs in the spot market in t .

The buyer, g , observes its time-specific ϕ_{gt}^s as well as all w_{gt} demand shocks before deciding whether to request a contract shipment in t , buy from the local spot market, or not receive any shipments.

²³ S_{lgt}^c is indirectly decreasing in all other demand shocks, which increase the number of contract requests and, hence, the fixed cost for each contract shipment fulfilled by l in t .

²⁴Appendix B describes how we measure spot market thickness.

Seller l 's payoff from selling a given shipment in the spot market is highest spot price available at time t , $\max_G(p_{gt}^s)$.

5.4 Equilibrium

Whenever a contract shipment is observed, g has found it incentive compatible to request that shipment in t , and l has found it incentive compatible to fulfil the request. Following the evidence in Section 3, showing that contract volumes increase at times when there is an adverse weather shock, we assume that g 's requests are consistent with per-period incentive compatibility.

To explain the fact that sellers choose to fulfil contract shipment requests even in times when their spot market outside alternative is particularly valuable, we assume the contractual relationship is dynamically incentive compatible for sellers because they extract value from the continuation of the relationship. Failing to deliver, on request, a contract shipment that fits buyer needs would end the relationship. That is, buyers play a trigger strategy in equilibrium.

5.4.1 Buyers' incentive compatibility constraints

The state variable of buyer g in t is its pre-determined set of contract partners, the demand shock, w_{gt} , and the realized spot market value shock ϕ_{gt} . Having also observed all other demand shocks, g forms an expectation of the local spot price, $E(p_{gt}^s)$.

Requesting a contract shipment is incentive compatible for g in t iff:

$$(1 - \beta)S_{lgt}^c \geq V_{gt}^s - E(p_{gt}^s). \quad (12)$$

From equations (10) and (11), this inequality is satisfied when ϕ_{gt} exceeds ϕ_{lgt}^{s*} , which is given by:

$$\phi_{lgt}^{s*} = \frac{v(w_{gt})}{(1 - \beta) \left[v(w_{gt})^{\frac{1}{1-\mu}} \left(\frac{\beta\mu}{c} \right)^{\frac{\mu}{1-\mu}} - F_{lt} \right] + E(p_{gt}^s)}. \quad (13)$$

Intuitively, g calls in a contract shipment from contract partner l when the demand shock is sufficiently high that g 's share of the contract surplus, given l 's relationship-specific effort, exceeds g 's surplus from the spot market, given the expected spot price and the realized ϕ_{gt}^s draw. The threshold value ϕ_{lgt}^{s*} varies with the demand shock because it is a function of contract shipment surplus.²⁵

²⁵Although we do not currently use it in the estimation, there is a second threshold value of ϕ_{gt}^s , above which the buyer prefers no shipment to a spot shipment, that is $\frac{v(w_{gt})}{\phi_{gt}^s} - E(p_{gt}^s) \leq 0$. This threshold value is:

$$\bar{\phi}_{gt}^s = \frac{v(w_{gt})}{E(p_{gt}^s)}. \quad (14)$$

If gt faces a positive demand shock that is sufficiently large that $S_{lgt}^c \geq 0$, the probability that g requests a contract shipment from l is the probability $\phi_{gt} \geq \phi_{lgt}^{s*}$, which is:

$$\Pr(\phi_{gt}^s \geq \phi_{lgt}^{s*}) = 1 - F_\phi(\phi_{lgt}^{s*}) = (\phi_{lgt}^{s*})^{-\lambda},$$

and the probability that g goes to the spot market is the probability that ϕ_{gt}^s is less than ϕ_{lgt}^{s*} . This probability is given by:

$$\Pr(\phi_{gt}^s \leq \phi_{lgt}^{s*}) = F_\phi(\phi_{lgt}^{s*}) = 1 - (\phi_{lgt}^{s*})^{-\lambda},$$

since whenever $S_{lgt}^c \geq 0$, the probability there are no shipments is zero.

5.4.2 Sellers' incentive compatibility constraint

We note that g is more likely to request a shipment when they receive a more positive demand shock, which is exactly when l 's outside option of selling on the spot market is likely to be of high value (since demand shocks are correlated within each region containing multiple spot markets). We impose that the seller's equilibrium strategy satisfies their dynamic incentive compatibility constraint. That is, l fulfils a contract request if, over the course of all future periods ($t' > t$), the share of contract surplus it will receive, β , less the variable effort costs it will incur, outweighs the revenue it would earn from selling the same quantities in the spot market with the highest spot price in each period $t' \geq t$.

Seller l 's incentive compatibility constraint for all future periods $t' > t$ can be written as follows:

$$\sum_{t' \geq t} \delta_{lt'} (\beta S_{lgt'}^c - c(1 + h_{lgt'})) \geq \sum_{t' \geq t} \delta_{lt'} E(\max_{g \in G} (p_{gt'}^s)), \quad (15)$$

where $g \in G$ indexes all destinations with spot markets. The discount factor is $\delta_{lt'} = \frac{1}{1+r_{lt'}}$, where $r_{lt'}$ is the seller-time-specific real interest rate.²⁶

6 Estimation

We now describe how we estimate the model set out in Section 5. We derive theoretical moments that govern whether g requests a contract shipment in timer t that is fulfilled by l . We then discuss the subsample of data used for estimation, which is chosen to reflect some of the abstractions made in the model. Finally, we describe the empirical moments and the estimation procedure.

6.1 Theoretical moments

The parameters of interest are those that relate the demand shock to shipment value (in the function $v(w_{lgt})$), and to endogenous effort via effort intensity, μ , the share of seller surplus, β , and the marginal

²⁶These country-year interest rates are taken from the World Development Indicators.

costs of contract effort, c . The other parameters relevant to contract shipments are those that affect the fixed costs, F_{lt} . We also focus on how the shape parameter of the Pareto distribution of spot value shocks, λ , relates to spot market thickness.

We specify the shipment value function to be a power function, governed by two parameters, α_0 and α_1 , so $v(w_{lgt}) = \alpha_0 w_{lgt}^{\alpha_1}$. Because the relationship between spot market thickness and its value will be of key importance in the counterfactual analysis, we estimate λ flexibly to be $\lambda = \lambda_0 (Th_{gt})^{\lambda_1}$, where Th_{gt} is the measure of spot market thickness. In order to allow the fixed costs of fulfilling a contract shipment to vary with deviations from the typical number of l 's contract shipments, we specify $F = F_0 (N_{lt})^{F_1}$, where N_{lt} is the deviation from the mean number of contract shipments a seller fulfils in each period.

Considering also effort intensity, μ , the variable cost of seller effort, c , as well as the share of contract surplus going to the seller, β , we have a total of nine parameters to estimate, $\theta = \{\alpha_0, \alpha_1, \lambda_0, \lambda_1, \mu, c, F_0, F_1, \beta\}$.

We first look at the implications of g 's incentive compatibility constraint in inequality (12). The threshold value of the shock above which a contract shipment is preferred to a spot shipment, ϕ_{lgt}^{s*} , from equation (13), can now be written:

$$\phi_{lgt}^{s*} \equiv \frac{\alpha_0 w_{lgt}^{\alpha_1}}{(1 - \beta) S_{lgt}^c + E(p_{gt}^s)}.$$

Hence, the theoretical moments giving the probability a contract shipment is observed in lgt , conditional on $S_{lgt}^c \geq 0$, are:

$$\Pr(\phi_{gt}^s \geq \phi_{lgt}^{s*}) = \begin{cases} \left(\frac{\alpha_0 w_{lgt}^{\alpha_1}}{(1 - \beta) S_{lgt}^c + E(p_{gt}^s)} \right)^{-\lambda} & S_{lgt}^c \geq 0 \\ 0 & S_{lgt}^c < 0 \end{cases} \quad (16)$$

where the term S_{lgt}^c is the surplus from a contract shipment and, under the functional form assumptions, equation (10) becomes:

$$S_{lgt}^c = (\alpha_0 w_{lgt}^{\alpha_1})^{\frac{1}{1-\mu}} \left(\frac{\beta \mu}{c} \right)^{\frac{\mu}{1-\mu}} - F_{lt},$$

and the theoretical moments in expression (16) are functions of all the model parameters of interest.

Some intuition for parameter identification can be found by observing that if w_{gt} is held constant, then the numerator and first term in the denominator of the theoretical moments are fixed, and variation in the probability of observing a contract shipment is determined only by variation in expected spot prices, allowing us to identify λ . Similarly, given the functional form assumptions for V_{gt}^o , $o \in \{c, s\}$, and holding $E(p_{gt}^s)$ fixed, variation in the probability of observing a contract shipment comes only from variation in w_{gt} , allowing us to estimate k, α, μ and a term that is a function of both β and c .

We then use l 's dynamic incentive compatibility constraint, inequality (15), to derive lower bounds for β and permit an estimate of c .

6.2 Data selection for model estimation

Because the model in Section 5 is at the lgt level, it abstracts away from several key aspects of the buyer and seller’s broader contexts. For the purposes of estimation, we select a subset of the available data to limit the impact of this simplification.

First, we restrict attention to lg routes where there is a contract in place.²⁷ For the seller, the route-week model we use does not account for the fact that the total quantity a seller can dispatch in any one week is more or less fixed. To ensure these seller constraints do not introduce biases into our estimation, we focus only on lgt observations where the seller sells positive quantities in at least one spot market, meaning they could have sold more via contract if they had received further requests that they chose to fulfil.

We also only include lgt observations where the buyer received at least one contract shipment and/or spot shipment from some seller $g \in G$. This ensures that $S_{lgt}^c \geq 0$. Hence, the buyer’s decision is only whether or not to request a contract shipment, conditional on having positive demand.²⁸

There are two further gt -specific variables that we construct to use as model inputs. The first is the number of contract shipments l fulfils in t relative to the typical number, N_{lt} , which helps us understand how sellers’ fixed costs of contract shipments vary with demand. N_{lt} is defined as the de-measured and normalized number of contract shipment requests that the seller delivers that month, where the number is de-measured by the mean number of contract requests l fulfils and the normalization ensures $N_{lt} \in (0, 1)$. Scheduling vessel capacity and loading slots to fulfil specific contract requests is an ongoing seller-level management challenge and the more requests they receive in the weeks around the week in question, the more costly it could be to allocate resources to any one individual contract request. We use deviations from the typical share of contracts to allow for the possibility that sellers who routinely handle a large number of contract shipments have likely made sunk investments in this fixed managerial capacity.

The second is the measure of spot market thickness, Th_{gt} , that allows us to estimate how the idiosyncratic shock to spot shipment value varies with the nature of the local spot market. Our preferred measure has two components. We consider the relative size of the spot market in the region in that period compared to the destination’s contract demand in that period. If buyers have total demand for a large number of shipments, and much of it is fulfilled via contract, then their spot market options are more limited. We make an adjustment to this measure by multiplying it by the extent of overlap between the distribution of vessel capacities requested by the buyer under contract and the distribution of vessel capacities available

²⁷These are the route-years where we see at least one contract shipment on the route that year.

²⁸Our model suggests weeks where buyers receive both contract and spot shipments are times when the buyer would have preferred to receive more contract shipments, but their additional requests were unfulfilled by the seller(s). We implicitly assume that, when they are asked, sellers $l \in L$ send sufficient contract shipments on a route to avoid relationship breakdown.

on the region-specific spot market. This adjustment accounts for the fact that even if the regional spot market is relatively large, a buyer might not be able to find a vessel that fits their demand specifications at any point in time.²⁹

6.3 Empirical moments and estimation

There are 23,541 lgt observations with positive trade flows that satisfy the restrictions in subsection 6.2. We construct an indicator, C_{lgt} equal to 1 if there is at least one contract shipment in lgt . This indicator is equal to 0 if there are no contract shipments, which, given sample restrictions, implies g buys on the local spot market in t . $C_{lgt} = 1$ for 13,689 of the 23,541 observations, which is 58%.

Because l is defined at the country level, other than for Australia, Russia, and the United States, and g is also at the country level, there are many lgt observations where there is more than one contract shipment and/or g buys more than one spot shipment. In the estimation, we weight observations by the sum of the number of spot shipments g buys in t plus the number of contract shipments in lgt . There are 60,287 weighted observations, of which $C_{lgt} = 1$ for 50,327, or 83%.

As in Section 3, we use weather shocks as the demand shocks w_{gt} . In the structural estimation, the weather shock is defined as the deviation from the average mean temperature in a destination in a given month so, for example, it tells us when January is particularly cold relative to a typical January in g .³⁰ Although local spot price variation is associated with weather shocks, there is sufficient independent variation in the data that both spot price and weather shocks are significantly correlated with the probability of receiving a contract shipment when both are included in a regression. We will use the coefficients in this regression as untargeted moments in our estimation as a robustness check for the parameter estimates. The weighted regression generating these two coefficients from the data sample is:

$$C_{lgt} = \alpha_1 w_{gt} + \alpha_2 p_{gt}^s + \gamma_{gm} + \varepsilon_{lgt} \quad (17)$$

The estimated coefficients are $\hat{\alpha}_1 = 0.069$ and $\hat{\alpha}_2 = 0.074$, and both are significant at the 5% level.

We form the joint log likelihood of weeks where contract shipments are observed by taking the product across all lgt :

$$L = \sum_{lgt} [C_{lgt} \log \Pr(\phi_{gt}^s \geq \phi_{lgt}^{s*}) + (1 - C_{lgt}) \log(1 - \Pr(\phi_{gt}^s \geq \phi_{lgt}^{s*}))], \quad (18)$$

Substituting for the theoretical moments giving the probability a contract shipment is observed in lgt , from

²⁹Appendix B gives a detailed description of how Th_{gt} is constructed.

³⁰Work in progress estimates the model using the standardized weather shock described in Section 2.3 rather than simply the deviation from the destination-month mean temperature. Comparing the reduced form results suggests the results will be qualitatively similar.

expression (16), and one minus this probability for the probability a spot shipment is observed, gives:

$$L = \sum_{lgt} \left[C_{lgt} \log \left(\frac{\alpha_0 w_{gt}^{\alpha_1}}{(1-\beta)S_{lgt}^c + E(p_{gt}^s)} \right)^{-\lambda} + (1 - C_{lgt}) \log \left(1 - \left(\frac{\alpha_0 w_{gt}^{\alpha_1}}{(1-\beta)S_{lgt}^c + E(p_{gt}^s)} \right)^{-\lambda} \right) \right], \quad (19)$$

where spot prices are the gt -level data shown in Figure 1.³¹ We estimate the model by maximizing the log likelihood in equation (19).³²

7 Results

7.1 Parameter estimates

Estimating the model by maximizing the log likelihood in equation (19) gives the following parameter estimates:

Parameter	Description	Estimate	Confidence Interval
α_0	Value function normalization	9.95	[9.83, 10.07]
α_1	Value function elasticity	0.21	[0.17, 0.25]
μ	Effort intensity	0.45	[0.44, 0.45]
c	Marginal effort cost	1.11	[1.10, 1.13]
F_0	Contract fixed cost normalization	4.98	[2.53, 7.42]
F_1	Contract fixed cost elasticity	0.61	[0.31, 0.84]
λ_0	Spot market thickness normalization	0.79	[0.69, 0.97]
λ_1	Spot market thickness elasticity	0.33	[0.29, 0.37]

All estimates are significantly different from zero. The positive value function parameters α_0 and α_1 show that the value g receives from a shipment is higher when there is a weather shock, that is, when it is unexpectedly cold. The positive estimates of λ_0 and λ_1 show that the value g receives from spot market shipments is positively associated with Th_{gt} , our measure of local spot market thickness at time t . The coefficient μ of 0.45 shows that relationship-specific seller effort contributes positively to contract shipment value. The positive coefficients F_0 and F_1 suggest there are fixed costs associated with each contract shipment that are increasing in the deviation from the typical share of contract requests that a seller fulfills that month.

³¹Rather than use spot price per cubic metre, we have used spot price per mmbtu in the estimation, which is the unit more often referenced. We normalize prices, so the unit of analysis does not affect the findings.


³²The estimation delivers estimates of confidence intervals for each parameter by computing the diagonal elements of the inverse of the Fisher Information Matrix, which is the negative Hessian found via numerical methods to approximate the second-order partial derivatives.

To find the β parameter estimate, we return to inequality (15):

$$\sum_{t' \geq t} \delta_{lt'} (\beta S_{lgt'}^c - c(1 + h_{lgt'})) \geq \sum_{t' \geq t} \delta_{lt'} E(\max_{g \in G} (p_{gt'}^s)),$$

which delivers an estimate of β for each lgt . The distribution of these estimates is given in Figure 5. Since

Figure 5: Estimated lower bounds on β



we need our estimate of β to satisfy all of these constraints, we take it as the maximum of the distribution in Figure 5, which is 0.2880, with a confidence interval [0.2846, 0.2915].

The results give that the seller exerts relationship-specific effort, that is, $h > 0$, for 49% of contract shipments. The seller incurs additional variable costs in these cases. This suggests that the seller's share of the contractual relationship surplus is sufficient to compensate them for foregoing spot markets and also exerting effort for these shipments.³³ In the model, h is an increasing function of the weather shock, and there is a strong positive correlation between estimated h_{lgt} values and w_{gt} . That is, sellers exert effort within contractual relationships when their contract partners face unexpectedly cold weather.

The parameter predictions allow us to predict the responsiveness of C_{lgt} , whether or not a contract shipment is observed in a route-week as determined by the estimated model, to observed variation in weather and spot prices using equation (17). These untargeted moments are close to the actual data, with the model predicting slightly more variation with spot prices than we see in the data.

7.2 Shipment value estimates


The estimated values of the observed contract shipments are shown in Figure 6. The far right histogram presents calculated shipment values from equation (6):

$$V_{gt}^c(w_{gt}, h_{lgt}) = \alpha_0 w_{gt}^{\alpha_1} (1 + h_{lgt})^\mu$$

³³We note that the seller receives only a β share of the additional surplus generated through its effort, and so is suboptimal.

Coefficient	Data	Model predictions
w_{gt}	0.069** (0.028)	0.070*** (0.001)
p_{gt}^s	0.074** (0.012)	0.100*** (0.001)

Figure 6: Estimate value of contract shipments




The middle histogram subtracts estimated fixed costs, F_{lt} , giving the distribution of S_{lgt}^c . The left (purple) histogram is the surplus of contract shipments net of estimated variable seller effort costs, $c(1 + h_{lgt})$. The mean of the left(purple) distribution is 5.58 and the standard deviation is 0.67. These values can be compared to the normalized distribution of the local spot price per mmbtu, the mean of which is 0.31. That is, we estimate that the mean value of a contract shipment is eighteen times higher than the mean spot price over the period studied.

The expected spot market shipment value can be calculated from local weather and the estimate of $\lambda_{gt} = \lambda_0 Th_{gt}^{\lambda_1}$, where Th_{gt} is the observed thickness of the local spot market.³⁴ It is:

$$E(V_{gt}^s) = E\left(\frac{\alpha_0 w_{gt}^{\alpha_1}}{\phi_{gt}^s}\right) = \alpha_0 w_{gt}^{\alpha_1} \frac{\lambda}{\lambda + 1} \quad (20)$$

Figure 7 shows the calculated expected spot values together with the net contract surplus net of variable effort costs (purple histogram) from Figure 6. In the figure, the blue histogram gives the expected values

Figure 7: Expected value of spot and contract shipments



of spot shipments when at least one spot shipment is observed in lgt , weighted by the actual number of spot shipments observed in the data. The mean of the distribution is 2.19 and the standard deviation is 0.41. The estimated mean value of a spot shipment is hence only 40% of the mean value of a contract shipment, but is still around 7 times larger than the normalized mean spot price. In other words, we find that around 15% of the gains from trade in the current spot market go to the seller and around 85% goes to buyers. Note that we estimate $\beta = 0.29$, which is the share of surplus going to sellers in the contract

³⁴The derivation is: $E(V_{gt}^s) = E\left(\frac{\alpha_0 w_{gt}^{\alpha_1}}{\phi_{gt}^s}\right) = \alpha_0 w_{gt}^{\alpha_1} \frac{\lambda_{gt}}{\lambda_{gt}+1} = \alpha_0 w_{gt}^{\alpha_1} E\left(\frac{1}{\phi_{gt}^s}\right) = \alpha_0 w_{gt}^{\alpha_1} \int \frac{1}{\phi_{gt}^s} \frac{\lambda_{gt}}{(\phi_{gt}^s)^{\lambda_{gt}+1}} d\phi_{gt}^s = \alpha_0 w_{gt}^{\alpha_1} \lambda_{gt} \int \frac{1}{(\phi_{gt}^s)^{\lambda_{gt}+2}} d\phi_{gt}^s = -\alpha_0 w_{gt}^{\alpha_1} \frac{\lambda_{gt}}{\lambda_{gt}+1} (\phi_{gt}^s)^{-(\lambda_{gt}+1)} \Big|_1^\infty = \alpha_0 w_{gt}^{\alpha_1} \frac{\lambda_{gt}}{\lambda_{gt}+1}$.

relationships, suggesting that the sellers extract a smaller share of the smaller gains from trade in spot shipments.

8 Counterfactual Analysis

We now ask how much long-term contractual relationships contribute to industry value relative to a situation where the same vessels arrived in each destination in each period and were transacted in the spot market. That is, the counterfactual is larger spot markets in each lgt where contract shipments are currently observed. The trade-off is that all shipments become “anonymous” arm’s length transaction, and sellers exert no relationship-specific effort and incur zero contract fixed and variable costs.

We are able to estimate this counterfactual because the λ parameter estimates from Section 7 allow us to calculate the expected spot shipment value as a function of the observed weather shock and the counterfactual spot market thickness, Th_{gt} , using equation (20). The measure of Th_{gt} typically becomes equal to 1 because the share of regional supply that is spot is now 100%. In addition, because the vessel sizes that were demanded under contract by each buyer are now a subset of the distribution of vessel sizes available via spot, the distributions of vessel sizes become equivalent (see Appendix B).

One major appeal of this approach is that we do not need to simulate counterfactual spot prices. We also avoid needing to address the seller’s allocation problem, that is, where they would send vessels in the absence of contracts. A related benefit is that our counterfactual estimates are invariant to transport costs, which are currently excluded from the analysis, but which remain constant in the counterfactual.

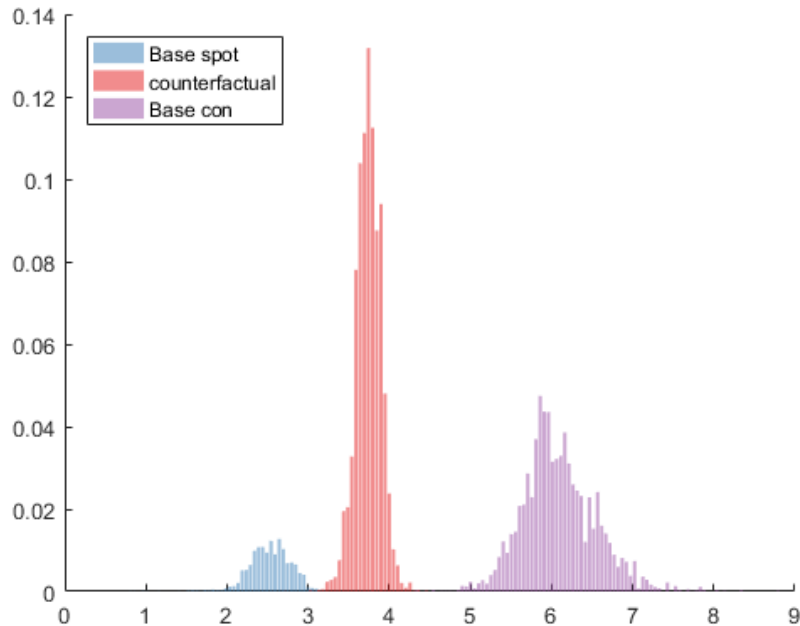
For this analysis, we estimate the counterfactual values of all contract and spot shipments, including those that did not satisfy the estimation sample restriction in Subsection 6.2. This gives us 31,428 contract shipments and 6,201 spot shipments, so 37,449 shipments in total.

The predicted values of these 37,449 shipments sold in the enlarged spot markets is given in the red histogram in Figure 8. For comparison, the blue histogram is the distribution of actual spot shipment values given the observed much smaller number of spot shipments, as shown in Figure 7. Figure 8 also includes the observed shipment surplus for contracts net of variable costs (the purple histogram). The red histogram reflects higher counterfactual spot values than in the blue histogram because buyers expect to receive a spot shipment that is a closer fit to their needs in larger spot markets. This counterfactual distribution has a mean value of 3.41 and a standard deviation of at 0.35, which is smaller than the standard deviation in the observed spot value distribution. This is because there is no variation in Th_{gt} in the counterfactual.

Because the counterfactual holds all transport costs fixed, and varies only the value created for buyers, it is possible to compare total welfare in the observed and counterfactual cases.

The total change in welfare can be decomposed into three parts, where the variables with asterisks

Figure 8: Observed and Counterfactual Shipment Values



denote the number of counterfactual shipments, Q_s^* , and the expected value of the distribution of the expected value of those counterfactual shipments, $E(V_s^*)$:

$$\begin{aligned}
 \Delta W &= \Delta W_s + \Delta W_c \\
 &= Q_s^* E(V_s^*) - (Q_s E(V_s) + Q_c (V(h) - F - c(1 + h))) \\
 &= (Q_s + Q_c) E(V_s^*) - (Q_s E(V_s) + Q_c (V(h) - F - c(1 + h))) \\
 &= Q_s (E(V_s^*) - E(V_s)) + Q_c (E(V_s^*) - V(h)) + Q_c (F + c(1 + h))
 \end{aligned} \tag{21}$$

The first term in equation (21) is the percentage increase in the value of spot markets in the counterfactual where they are much larger. This change is 3.43%. The second term is the percentage reduction in the value of the shipments that are currently transacted via contracts, and is estimated to be 73.65%. This loss can be attributed to the loss in shipment value from the fact that sellers no longer match these shipments to buyers' exact needs and exert no relationship-specific costly effort to ensure the quality of the match. The third term in equation (21) is the welfare gain coming from the savings in the seller's variable effort costs and fixed costs associated with contract shipments. This saving is equivalent to 38.06% of industry value over the time period studied.

The net impact of the removal of contracts is a reduction in industry value of 32.12%. We infer that total industry size remains too small for spot markets to deliver shipments to buyers that fit their demand even if all industry supply were transacted on global spot markets rather than via contracts.

9 Conclusion

Long-term contracts dominate the LNG industry and the nature of these relationships shapes industry value. This paper brings insights from recent advances in the theory of incomplete contracts to bear on detailed transaction-level data. Our main contribution is the finding that these contracts allow transactions to flexibly adjust to unexpected demand circumstances. We also quantify the significant role played by relationship-specific seller effort in creating value for the buyer. Adaptation, often enabled by non-contractible effort, increases industry efficiency relative to a counterfactual of LNG spot markets characterized by frictions from technology and geography. The analysis reveals that the value of the global LNG industry would only have been two thirds of its actual size between 2009 and 2019 in the absence of adaptation within long-term contracts.


Appendix

A Global supply and demand in 2019

A.1 Supply

The shipment-level data we use, which runs from the start of 2009 to the end of 2019, contains shipments from 19 countries. Qatar has been the largest exporter historically, with over 500 shipments in each of the 11 years. Australia has increased exports over the time period, and exported more than 500 shipments in the last four years in the data. After Qatar and Australia, the largest exporter countries are Malaysia, Indonesia, Algeria, and Nigeria. Figure A.1 focuses on 2019, and shows total volumes exported by source country in cubic metres. The USA, Russia, and Trinidad and Tobago are more prominent in 2019 than in the period as a whole.


Figure A.1: Largest exporter countries in 2019



Because the LNG industry is a costly shipping technology for a commodity product, we might think that both sellers and buyers have incentives to minimize the distance travelled. The overall flows confirm this, in general. For example, Figure A.2 shows that the two largest destinations for Australia are the relatively close Japan and China, where Australia supplies 39% and 47% of these countries' LNG, respectively. Qatar supplies less of these countries' demand, at 11% and 13%. The figures below the country name show the average mileage travelled on routes between the source and destination countries. The LNG travelling to Japan and China from Australia travels much shorter distance than the LNG from Qatar.

However, there are exceptions. For example, Qatar supplies 27% of South Korea's LNG and 29% of

Figure A.2: Volume by Destination, 2019



Taiwan’s despite LNG having to travel over 5,000 miles to get there. Australia is much closer to both destinations and supplies only 19% and 27%, respectively.³⁵ We suspect deviations from gravity relate to the staggered timing at which liquefaction and regasification capacity was installed in different countries and the agreed long term contracts between buyers and sellers.

Figure A.2 also offers some insight about how shipments are organized. The bulk of LNG exports from Qatar in 2019 were contracted quantities, meaning very little LNG from Qatar ended up in the spot market. For Australia, spot volumes were relatively more common.

Figure A.3 focuses on 2019, and shows total volumes imported by destination country.

Figure A.4 shows the volumes arriving in Japan in 2019 by source country. The figures below the country names are the average distance travelled and the share of that country’s LNG exports arriving in Japan. There are some interesting comparisons. Australia sends Japan 39% of its output but is much further away than Malaysia, which sends a similar share of its total output to Japan.

B Measuring spot market thickness, Th_{gt}

. For each destination-week, the measure Th_{gt} is the product of two terms.

First, we find the share of spot shipments in the region (Europe or Asia) in that month, which is the total number of spot shipments delivered to all destinations in the region divided by the total number of spot and contract shipments delivered in the region. The higher is this share, the more likely any buyer in

³⁵We have some data on production costs. Qatar has a much lower operating production cost than the installations in Australia (see supply functions in earlier update), so the delivered cost from Qatar may still be much lower than from Australia.

Figure A.3: Largest importer countries in 2019


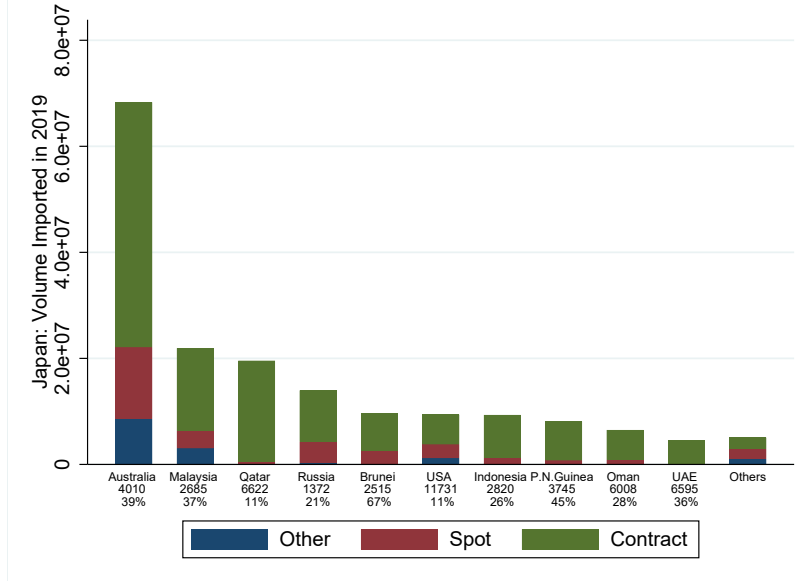


Figure A.4: Japan imported volumes by Origin, 2019



the region could purchase spot shipments at the time they demand it.

Second, we turn to the match of spot shipment capacities with those demanded via contracts. We find the distribution of spot vessel sizes that have been delivered in the region over the last three months. We then add to these shipments the contract shipments delivered to the destination that month to give the capacity distribution of typical region spot shipments together with the destination's contract capacity demand. We use the Kolmogorov-Smirnov statistic to tell us how similar these two distributions are. If they are very similar, we infer that the destination's demand contract capacities could be catered to by the region's spot market. The more similar the distributions the thicker the local spot market, in terms of shipment capacity availability.

In the counterfactual, by construction, both of these terms equal one. Hence, the spot market in each destination-week achieves the maximum spot market thickness that is ever observed in the actual data. The buyer values for this enlarged spot market are still discounted due to the value of ϕ_{gt}^s in V_{gt}^s being determined by the estimated λ_0 and λ_1 parameters delivered in Section 7.

References

- Brancaccio, Giulia, Myrto Kalouptsi, and Theodore Papageorgiou. 2020. "Geography, transportation, and endogenous trade costs." *Econometrica* 88 (2):657–691.
- Costinot, Arnaud, Lindsay Oldenski, and James Rauch. 2011. "Adaptation and the boundary of multinational firms." *The Review of Economics and Statistics* 93 (1):298–308.
- Forbes, Silke J and Mara Lederman. 2010. "Does vertical integration affect firm performance? Evidence from the airline industry." *The RAND Journal of Economics* 41 (4):765–790.
- Forbes, Silke Januszewski and Mara Lederman. 2009. "Adaptation and vertical integration in the airline industry." *American Economic Review* 99 (5):1831–1849.
- Grossman, Gene M, Elhanan Helpman, and Hugo Lhuillier. 2023. "Supply Chain Resilience: Should Policy Promote International Diversification or Reshoring?" *Journal of Political Economy* forthcoming.
- Grossman, Gene M, Elhanan Helpman, and Alejandro Sabal. 2023. "Resilience in Vertical Supply Chains." Tech. rep., National Bureau of Economic Research.
- Harris, Adam and Thi Mai Anh Nguyen. 2022. "Long-term Relationships and the Spot Market: Evidence from US Trucking." .
- Joskow, Paul L. 1988. "Asset specificity and the structure of vertical relationships: empirical evidence." *The Journal of Law, Economics, and Organization* 4 (1):95–117.
- Klein, Benjamin, Robert G Crawford, and Armen A Alchian. 1978. "Vertical integration, appropriable rents, and the competitive contracting process." *The journal of Law and Economics* 21 (2):297–326.
- Kranton, Rachel E. 1996. "Reciprocal exchange: a self-sustaining system." *The American Economic Review* :830–851.
- Lafontaine, Francine and Margaret Slade. 2012. "Inter-firm contracts: Evidence." *Handbook of organizational economics* :958–1013.
- Macchiavello, Rocco. 2022. "Relational contracts and development." *Annual Review of Economics* 14:337–362.
- Macchiavello, Rocco and Ameet Morjaria. 2015. "The value of relationships: evidence from a supply shock to Kenyan rose exports." *American Economic Review* 105 (9):2911–2945.
- Masten, Scott E and Keith J Crocker. 1985. "Efficient adaptation in long-term contracts: Take-or-pay provisions for natural gas." *The American Economic Review* 75 (5):1083–1093.
- Nicholson, Sharon E. 1993. "An overview of African rainfall fluctuations of the last decade." *Journal of climate* :1463–1466.
- Pirrong, Stephen Craig. 1993. "Contracting practices in bulk shipping markets: a transactions cost explanation." *The Journal of Law and Economics* 36 (2):937–976.

- S&P. 2022. “Gunvor backs out of LNG deliveries to Pakistan for April-June deliveries.” <https://www.spglobal.com/commodityinsights/en/market-insights/latest-news/lng/032822-gunvor-backs-out-of-lng-deliveries-to-pakistan-from-april-june-officials>. Accessed: 2023-09-12.
- . 2023. “Demand for term LNG contracts firms amid supply security concerns.” <https://www.spglobal.com/commodityinsights/en/market-insights/latest-news/lng/071223-demand-for-term-lng-contracts-firms-amid-supply-security-concerns>. Accessed: 2023-09-12.
- Tolvanen, Juha, Olivier Darmouni, and Simon Essig Aberg. 2021. “Pulp Friction: The Value of Quantity Contracts in Decentralized Markets.” *Available at SSRN 4381156* .
- Williamson, Oliver E. 1979. “Transaction-cost economics: the governance of contractual relations.” *The journal of Law and Economics* 22 (2):233–261.
- Zahur, Nahim Bin. 2022. “Long-term contracts and efficiency in the liquefied natural gas industry.” *Available at SSRN 4222408* .