Queen's Economics Department Working Paper No. 1517

Market Power and Spatial Price Discrimination in the Liquefied Natural Gas Industry

Nahim Bin Zahur

Department of Economics
Queen's University
94 University Avenue
Kingston, Ontario, Canada
K7L 3N6

2-2023

Market Power and Spatial Price Discrimination in the Liquefied Natural Gas Industry

Nahim Bin Zahur *

February 2023

Abstract

The liquefied natural gas (LNG) industry is characterized by systematic inter-regional price differentials, raising the question of whether sellers price discriminate. This paper measures market power in the LNG spot market and studies how market power influences pricing, trade and welfare. I develop a novel method for inferring market conduct that utilizes information on sellers' pricing and quantity decisions across multiple geographically segmented markets. My test for market conduct is based on the observation that sellers exercising market power engage in third-degree price discrimination, whereas sellers behaving competitively do not. Using data from 2006 to 2017 on spot market trade flows, spot prices, shipping costs and seller capacities, I estimate a structural model of LNG trade and pricing that incorporates spatial differentiation, capacity constraints and trade frictions and flexibly nests different models of seller market power. I find that seller decisions are consistent with a Cournot model and unlikely to be generated by a competitive model. The total deadweight loss from market power is estimated to be USD 12 billion, or about 4.5% of total revenue. I find that market power plays a key role in exacerbating inter-regional price differentials.

Keywords: Market Power, Price Discrimination, Conduct Parameter, Contracts, Liquefied Natural Gas

---

*Department of Economics, Queen's University (email: nz17@queensu.ca). I am grateful to Panle Jia Barwick, Jean-Francois Houde, Shanjun Li and Giulia Brancaccio for invaluable advice and guidance. I thank participants at CEA, WEAI, ESAM, Energy Research Insights for Decisionmaking Conference, Cornell IO Workshop, and Cornell Dyson Energy and Environmental Economics Workshop, for helpful comments. Generous funding for this research project was provided by the Alfred P. Sloan Foundation Pre-doctoral Fellowship on Energy Economics, awarded through the National Bureau of Economic Research (NBER).

1 Introduction

The modern history of the liquefied natural gas (LNG) industry has been characterized by periods when spot prices differ sharply across markets. A case in point is the Fukushima nuclear disaster in March 2011, which prompted a shutdown of Japan's nuclear reactors and a sharp increase in Japan's LNG demand. During the year preceding the Fukushima disaster, Asian LNG spot prices were $1/MMBtu higher than European spot prices. During the three years after the disaster, though, Asian prices were on average $4.75/MMBtu (or 40%) higher than European prices.

These large price gaps are puzzling in light of several facts about the LNG industry. LNG is essentially a homogeneous product, with only minor differences in the quality of LNG produced by different suppliers. The market has grown rapidly in the past fifteen years: the total volume of trade has more than doubled, and the number of countries importing LNG has tripled. LNG trade, historically carried out almost entirely through long-term contracts, has shown signs of becoming increasingly flexible, with spot and short-term trade accounting for 26% of total trade volume since 2010. The increased size and liquidity of the market suggest that gas markets around the world should have become more integrated, but the large differences in prices across regions indicate otherwise.

The higher prices paid by buyers in Asia compared to buyers in Europe for the same commodity has sparked a debate about whether LNG suppliers are engaging in price discrimination. It has been common for LNG buyers in Asia to refer to an "Asian price premium" (see, for instance, Koyama, 2012). In response to high Asian prices and the perceived premium paid by Asian buyers, there have been several attempts by Asian importing countries to form buyer's alliances.[1] Yet the existence of pricing differentials across regions is insufficient by itself to establish that sellers exercise market power. The industry generally operates at very high levels of capacity utilization, demand is fairly inelastic, and shipping costs are important. Under these conditions, prices may differ across markets even under competitive behavior. Finally, though the LNG industry is concentrated and arbitrage is costly, LNG suppliers often have strong ties with governments and it is by no means a given that they would exercise market power in a way consistent with standard oligopoly models. Whether or not there exists market power in LNG is therefore an empirical question.

In this paper, I ask whether there is evidence that sellers exercise market power in the LNG spot market. The goals of the paper are to measure market power in the LNG spot market, and study how market power influences pricing, trade and allocative efficiency. My method for inferring market conduct utilizes information on both the pricing and quantity decisions of sellers across multiple geographically segmented markets. I build on the observation that sellers exercising market power

[1] For example in March 2017, three major LNG buyers located in Japan, Korea and China signed an MoU to form a buyer's alliance. See https://www.reuters.com/article/us-asia-lng-markets-idUSKBN16U27X.

engage in third-degree price discrimination by charging different prices (net of transportation costs) to different buyers, whereas sellers behaving competitively charge the same prices to all buyers. Moreover, in the presence of capacity constraints, discriminatory behavior on the part of sellers will distort the allocation of goods across buyers.

The key challenge in identifying market power is that in the presence of capacity constraints and shipping costs, a competitive model can also generate price dispersion across markets. Consider an industry with geographically differentiated buyers and sellers. Even in the absence of capacity constraints, competitive prices paid by different buyers will differ due to shipping costs: a buyer located far away from sellers will generally pay a higher price than a buyer with a more favorable location. If in addition sellers are capacity constrained, the net price received by sellers (i.e. price minus shipping costs) will also not be equalized across destinations, because a capacity-constrained seller will not sell to all possible destinations. Thus without any quantity information, price differentials, whether measured at the buyer or seller level, cannot be directly interpreted as evidence of market power. It is the combination of price differentials and quantity allocation decisions that indicate potential violations of the competitive model. If the net price a seller receives differs within the set of buyers to which they sell positive quantities, the seller is unlikely to be acting competitively. Finally, the existence of unobserved trade costs and other frictions (such as unobserved trade costs) can also result in deviations from competitive behavior, and the analysis needs to distinguish between market power and other trade frictions.

To diagnose market power in the LNG spot market, therefore, I first begin by investigating whether the evolution of spot prices and trade flows is consistent with competitive behavior. I document evidence of persistent spot price differentials across regions: as discussed above, Asian LNG importers paid significantly higher prices than European buyers during a period ranging from 2011 to 2014, in the aftermath of the Fukushima nuclear disaster. Though competitive behavior can generate price differences across markets, the pattern of spot trade that emerged during this period is difficult to reconcile with competitive behavior. In particular, a number of exporters sold spot LNG cargoes to both Asia and Europe despite the fact that they would have received significantly higher prices (net of shipping costs) by increasing their sales to Asia.

Although the pattern of sellers selling to multiple markets at different net prices suggests deviations from perfectly competitive behavior, it does not answer the question of whether these deviations are the result of market power or some other trade frictions. I thus carry out a reduced-form test for whether sellers exercise market power that is designed to tease apart the role of market power from other frictions. The intuition behind the test is that under the hypothesis of market power, large sellers have a weaker incentive to increase sales to a market with high demand than small sellers, since by withholding sales from that market they can push up the price and increase their total revenue from that market. As such, large sellers will respond less sharply to price changes than small sellers. Under perfect competition subject to trade frictions, by contrast, sellers ought to be equally responsive to price changes irrespective of their total sales. I find that spot sales of small sellers are indeed more responsive to price differentials than the spot sales of large sellers, consistent with the existence of market power in the LNG spot market.

In order to quantify the extent of market power, I develop and estimate a structural model of LNG trade and pricing that incorporates spatial differentiation, capacity constraints, trade frictions, and seller market power. The model features a set of sellers and buyers that are located in geographically distinct locations. Capacity-constrained sellers produce a homogeneous good that they can sell to the buyers, but only by incurring a shipping cost that is increasing in the geographical distance between the seller and the buyer, as well as an idiosyncratic trade cost that reflects unobserved trade frictions. Each buyer has a downward-sloping curve with a stochastic intercept that reflects buyer-specific demand shocks. Trade takes place via both long-term contracts and a spot market. Contracted quantities are fixed in the short-run and cannot adjust in response to the demand shocks. In the spot market, by contrast, sellers observe demand shocks prior to making any decisions. The model nests two distinct modes of seller behavior. Under competitive behavior, sellers choose their quantities taking prices as given. Under Cournot competition, sellers choose their own spot quantities taking as given the spot quantities chosen by their rivals as well as all contracted quantities. A key feature of the model is that because of capacity constraints, the seller's quantity decisions across markets is interdependent.

I estimate this model using data on capacity, trade flows, spot prices and shipping costs in the LNG industry from 2006 to 2017. The dataset has the unique feature that I separately observe trade flows that take place under long-term contracts from trade flows that take place under short-term contracts and spot transactions. I first estimate demand curves for each buyer, using aggregate export capacity and electricity demand in other regions as instruments for the spot price. I then use the trade flow data, spot prices and shipping costs to estimate a set of parameters that characterize seller export behavior. A key goal of the estimation exercise is to estimate a market conduct parameter that allows me to test whether exporting countries act strategically in the spot market. I show that seller conduct can be identified from a first-order condition that relates the seller's spot sales across buyers with the prices it receives from each buyer net of shipping costs. I estimate the conduct parameter as well as other parameters characterizing trade costs and capacity constraints, and find that seller decisions are consistent with a Cournot model and unlikely to be generated by a competitive model.

I then use the estimated model to carry out a simple counter-factual exercise where sellers behave competitively on the spot market and do not price discriminate. I find that the deadweight loss from market power is USD 12 billion over the 12-year period considered in the study, or about 4.5% of total spot market revenue. Under perfect competition, inter-regional price differentials are 45% smaller than under Cournot competition. This suggests that the large price differentials observed in the LNG market are to a significant extent driven by sellers exercising market power on the spot market.

My paper contributes to a growing literature detecting market power and studying the welfare consequences of market power in energy industries. A rich literature analyzes market power in restructured electricity markets in California (Borenstein et al., 2002; Puller, 2007) and elsewhere (Wolak, 2000; Fabra and Toro, 2005; Sweeting, 2007; Hortacsu and Puller, 2008). The role of OPEC in the oil market has long been the subject of academic study, with a recent example being Asker et al. (2019). In contrast, there has been much less work on market power in the LNG industry. The only existing paper that I am aware of is Ritz (2014), who develops a theoretical model to illustrate how LNG exporters can price discriminate across markets and argues that market power can rationalize observed inter-regional price differentials. My paper, in contrast, utilizes panel data on LNG trade flows and spot prices to structurally estimate a model of the LNG market. This allows me to both measure the level of market power and analyze the extent to which observed price differentials can be explained by market power, while accounting for the role played by capacity constraints, shipping costs and trade frictions.

The method developed here for identifying firm conduct builds on an extensive literature in empirical industrial organization on measuring market power (Bresnahan, 1982, 1987; Genesove and Mullin, 1998; Nevo, 2001; Salvo, 2010; Miller and Weinberg, 2017). Typically this literature identifies various models of firm conduct from data on each firm's total production, prices and costs: firms exercising market power produce lower output and charge higher markups than firms choosing quantities or prices competitively. By contrast, I identify market conduct from observed third-degree price discrimination and the extent to which firms with market power split sales across destinations. My strategy for inferring seller conduct is thus similar to that employed by Dafny (2010), who shows that health insurers price discriminate across employers and argues that this behavior is inconsistent with competitive behavior. One key difference from Dafny (2010) is that capacity constraints and shipping costs are important in my setting, meaning that seller conduct needs to be inferred not just from price dispersion but also from quantity allocation decisions.

Finally, my paper is related to the empirical literature documenting price discrimination and studying the effects of price discrimination in a variety of settings (Villas-Boas, 2009; Grennan, 2013; Miller and Osborne, 2014; Li, 2015; Boik, 2017). The most closely related paper to mine in this literature is Miller and Osborne (2014), who estimate a structural model of the Portland cement industry that features price discrimination by spatially differentiated suppliers facing capacity constraints. Unlike Miller and Osborne (2014), who impose the assumption that sellers engage in Bertrand-Nash competition when estimating their structural model, my estimation strategy allows the researcher to be agnostic about the oligopoly model characterizing the industry and instead infer the level of competition from sellers' pricing and quantity allocation decisions across markets. The data required by this strategy, however, is more stringent than Miller and Osborne (2014), whose model can be estimated using only partially observed and aggregated prices and quantities. The estimation strategy developed here is therefore useful in settings where the researcher observes detailed data on trade flows and prices charged in different markets but does not have prior knowledge of the nature of oligopolistic competition.

The analysis in this paper focuses exclusively on the spot market in LNG, taking as given both long-term contracts and export capacity. In a follow-up paper that I am currently working on, I endogenize both long-term contracts and investment and study the link between long-term contracts and efficiency. A key rationale behind long-term contracts is that they facilitate efficient levels of irreversible investments made by sellers faced with hold-up risk (Williamson, 1979). This is likely to be very relevant in the LNG industry, where an exporter must construct an export terminal that typically costs several billion dollars and at least 4 years to build before they can engage in LNG exports. However contracts can result in short-run inefficiency since they do not respond flexibly in response to demand shocks. In ongoing research I study the trade-off between hold-up risk (which is improved by contracts) and short-run misallocation (which is worsened by the use of contracts). Whether contracts are welfare-improving in the long-run depends on which of these two effects dominates. To understand the linkage between contracts and long-run efficiency, the paper develops and estimates a structural model of investment, contracts and trade in the LNG industry. By simulating the path of investments and trade in a counterfactual scenario where firms face restrictions on the usage of long-term contracts, I can assess whether long-term contracts lead to the efficient outcome, or whether there are welfare gains from more flexible trading arrangements. Initial estimates suggest that the short-run misallocation from long-term contracts is substantial and the inflexibility of long-term contracts also plays a role in creating the price differentials across regions that are characteristic of the LNG industry.

The remainder of this paper is divided as follows. Section 2 discusses key institutional features of the LNG industry and describes the dataset I use for the analysis. Section 3 provides descriptive evidence for market power, including a reduced-form test for whether sellers exercise market power. Section 4 develops a model of trade and pricing in this industry. Section 5 describes estimation of the key parameters of the model that characterize demand, cost and conduct. Section 6 concludes.

2 Industry and Data

Liquefied natural gas (LNG) is natural gas that has been cooled into a liquid form so that it can be transported in specialized LNG tankers. For LNG trade to take place, both the exporting country and the importing country need to invest in specialized infrastructure. In the exporting country, a liquefaction terminal converts natural gas into LNG, and loads the LNG onto a tanker. The tanker then transports LNG to a receiving terminal (known as a regasification terminal) at the importing country, where the LNG is unloaded. Once the LNG has been unloaded, it is usually converted back into gaseous form and used for power generation and for heating. In some countries with more limited domestic pipeline infrastructure, LNG is loaded into trucks at the import terminal and then transported via truck to the location where the natural gas will be used, though this is still fairly uncommon and accounts for less than 5% of total LNG imported globally.[2]

Many countries with limited gas reserves rely on imported liquefied natural gas (LNG) to meet their natural gas needs. For countries such as Japan and Korea that do not have their own gas reserves and cannot import gas via pipelines, LNG provides their only source of natural gas. Other countries, such as China and Spain, import natural gas via both pipelines and LNG. The largest producer of LNG is Qatar, followed by Australia, Malaysia and Indonesia. Japan is the largest importer of LNG, followed by China, Korea, India, Taiwan and Spain.

The first commercial LNG export facility was constructed in Algeria, which began exporting LNG in 1964. The United States began LNG exports from the Kenai Peninsula in Alaska to Japan in 1969. The industry has grown considerably in size since then, and industry growth has been especially rapid during the last fifteen years. The number of LNG importing countries has increased from 14 (in 2004) to 40 (in 2017), while the number of LNG exporting countries has increased from 12 (in 2004) to 19 (in 2017). The total volume of LNG trade has almost doubled between 2004 and 2017.

A key institutional feature of the LNG industry is that the the majority of trade is carried out under long-term contracts signed between LNG suppliers and downstream buyers. Long-term contracts are often signed prior to capacity investments, in part because without contracts investors in this industry face a hold-up problem (Williamson, 1979): an investor is in a weaker bargaining position if they seek trading partners after they have already committed to making the investment, since the cost of the investment is already sunk. A typical long-term contract specifies the average annual contracted quantity to be sold by the seller to the buyer, the start and the end date, and a pricing formula that is used to determine the price under which trade takes place. Typically the price of contracted LNG is indexed to the price of some benchmark. The price of crude oil is the most common such benchmark, while for terminals operating out of the US, the price is typically indexed to the domestic price of natural gas in the Henry Hub market. In addition to these basic features, a contract usually includes several additional clauses. The contract may specify a "take-or-pay" share: if the take-or-pay share is X% and the quantity agreed is Q, then the buyer commits to taking at least $X\%$ of Q every year; if the buyer fails to take delivery of the $X\%$ of Q in a given year, they must nevertheless pay for that LNG. This ensures that the seller's minimum revenue

[2]LNG trucking has been most extensively used in China, which accounted for 74% of trucked LNG globally in 2017.

from the contract is the price multiplied with the take-or-pay quantity.A contract may include a destination restriction, which means that the LNG cargo must travel only to the destination that is specified on the contract. Finally, the contract may also include a diversion clause which specifies how the two parties split the surplus in the event that they decide to sell the cargo to a third party (which is known in industry parlance as a “diversion").

Figure 1 shows the distribution of contract durations for long-term contracts (defined for the purposes of this study as contracts that are longer than 4 years in duration). The majority of contracts are well over 10 years in length, with the modal contract length being around 20 years.

Long-term contracts do not provide the only way for parties to trade LNG, however. Parties can also trade LNG on the spot market, or using short-term contracts, and this is the focus of this study. A typical spot transaction is negotiated between the buyer and the seller three months prior to delivery date, and will involve the delivery of a single LNG cargo from the seller’s export terminal to the buyer’s import terminal.[3] Figure 2 shows that the share of spot and short-term trade has been rising over time, accounting for only 12% of trade in 2004 but 27% of trade in 2017. Finally, buyers can also “re-export" LNG to other buyers by purchasing LNG from one source and re-loading the LNG onto a new tanker. The share of re-exports increased from 0.15% in 2008 to a peak of 2.69% in 2014, but decreased to 1% of total trade by 2017.

Figure 3 shows the evolution of industry capacity and trade over time. It plots total liquefaction capacity (which is the capacity of LNG exporters), total regasification capacity (which is the capacity of LNG importers) and total LNG trade. The binding constraint on the volume of trade tends to be the available liquefaction capacity, and export capacity utilization tends to be high, ranging from between 80 to 90%. By contrast, there is generally a lot of excess regasification capacity, and this has grown in recent years with increasing entry of new importing countries. These patterns reflect the fact that liquefaction plants are typically more costly to build than regasification plants.

Data

The empirical analysis will draw on historical data on the global LNG market, which I have collected from various industry sources.

Data on liquefaction and regasification capacity I obtain data on investment and capacity data, for both the liquefaction terminals (owned by exporters) and regasification terminals (owned by importers), from the annual reports of the GIIGNL (The International Group of Liquefied Natural Gas Importers), whose members collectively own nearly every LNG import terminal in the world. The dataset covers all liquefaction and regasification terminals operational from 2004 to

[3]Each LNG tanker will carry “one cargo". The volume of LNG carried by a fully loaded LNG tanker is on average around 132,000 m^3, though it differs depending on the size of the tanker. There were a total of 5019 cargoes delivered in the year 2017.

Figure 1: Histogram of long-term contract durations

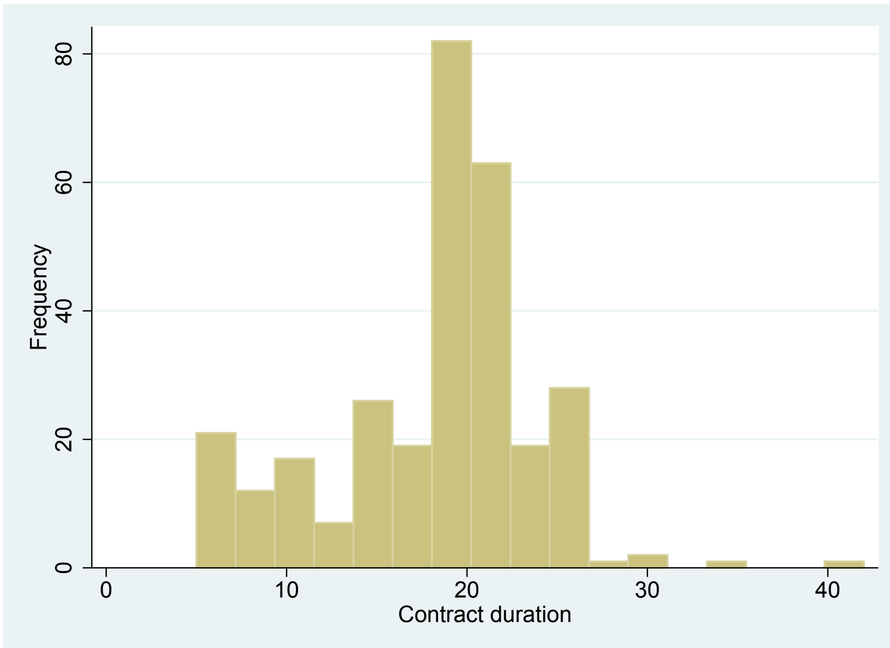

Note: long-term contracts are defined as contracts that are longer than 4 years in duration.

Figure 2: Share of short-term contracts and spot transactions in LNG trade

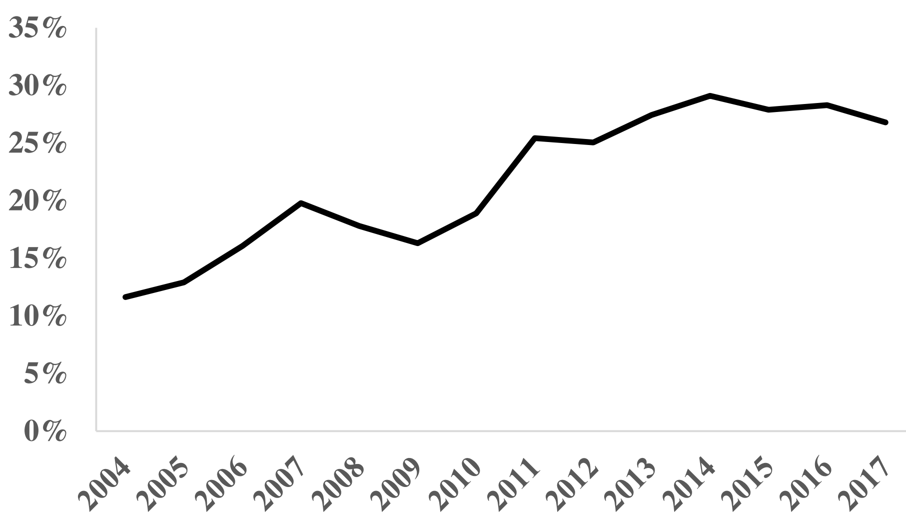

Note: the above figure plots trade carried out in the spot market or using short-term contracts as a share of total LNG trade. Short-term contracts are defined as contracts that are 4 years in length or shorter. Source: GIIGNL Annual Reports.

Figure 3: Liquefaction capacity, regasification capacity and LNG trade over time

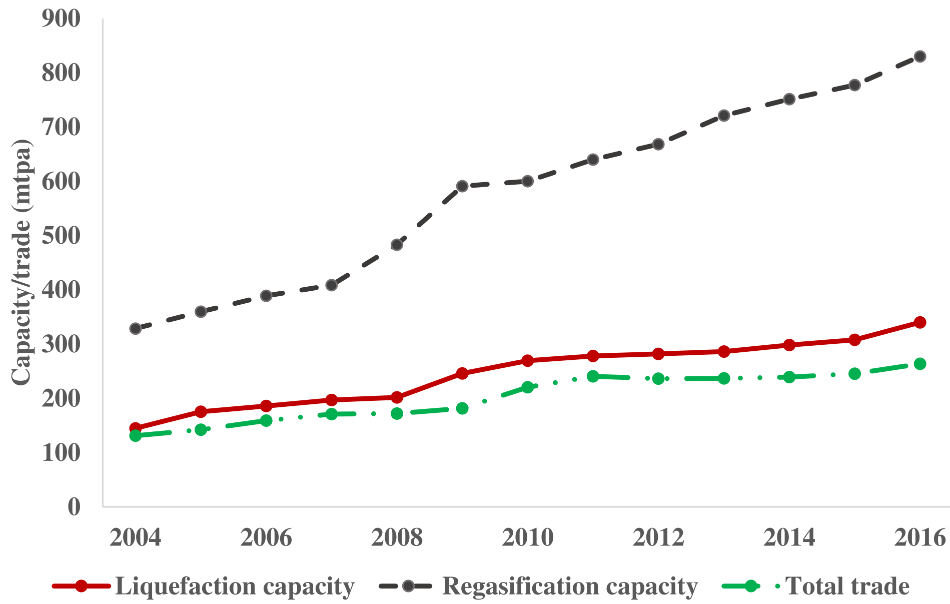

Note: The figure plots global liquefaction capacity, regasification capacity and the volume of LNG trade over time. Liquefaction capacity refers to the annual nameplate capacity of liquefaction plants (which are used to export LNG). Regasification capacity refers to the annual nameplate capacity of regasification plants (which are used to import LNG). Capacity is measured in million tons per annum (mtpa), while trade is measured in million tons (mt).

2017. The dataset consists of the start-up year, nameplate capacity, the ownership structure and the operator for every export and import terminal built on or prior to 2017. From 2004 onward, I also observe the year when a Final Investment Decision (FID) is made on an export terminal[4]. Due to time-to-build, the FID year is earlier than the start-up year (which is when the terminal begins operating).

Data on LNG trade flows I obtain data on LNG trade flows from the annual reports of the GIIGNL. This dataset reports yearly LNG trade flows for each country pair from 2004 to 2017. We separately observe trade flows that take place under long-term contracts (exceeding four years in length), short-term contracts and spot trade[5], and re-exports.[6] A novel feature of this dataset is that it allows me to distinguish between flows that take place under long-term contracts and short-term/spot contracts. Actual trade flows under long-term contracts may differ from the quantity originally agreed to in the contracts between the two countries, either due to production fluctuations or because the parties to the contract agreed to divert cargoes to other destinations. It should also be noted that this dataset record country-country trade flows, rather than firm-firm trade flows.

LNG spot prices and freight costs I obtain data on weekly LNG spot prices and shipping costs between February 2006 and August 2018. The most comprehensive of these datasets comes from Waterborne Energy[7]. The Waterborne Energy dataset reports weekly landed spot LNG prices (measured in USD/MMBtu) at 18 major LNG destinations, as well as weekly freight rates (in USD/MMBtu) for 220 exporter-importer pairs.[8]

I complement this dataset with additional spot price information from a number of sources. Thomson Reuters publishes a spot price index for North-east Asia covering the period from July 2011 to August 2018, as well as three indices for spot prices in Singapore, North Asia and Dubai/Kuwait/India (from October 2014 to August 2018) that is published by Singapore Exchange (SGX) and Energy Market Company (EMC). Finally, in the US and UK, spot LNG prices are closely related to the domestic price of natural gas. I obtain the Henry Hub natural gas price in the US from the Federal Reserve Economic Data (FRED) database, and the National Balancing Point (NBP) gas price series in the UK from Bloomberg.

Data on LNG contracts I also observe data on individual LNG contracts. This data is avail-

[4]Typically the Final Investment Decision (FID) is preceded by various feasibility studies, obtaining the relevant regulatory permits and signing sale and purchase agreements or contracts with buyers. The FID itself is the decision by all project partners to finally commit to the project. Construction of an LNG export terminal only begins once a Final Investment Decision has been taken by the investors.

[5]Short-term contracts are defined as contracts that are four year or shorter in duration. The dataset does not separately distinguish short-term contracts from one-time spot transactions.

[6]This records exports from one importing country to another importing country, which is distinct from exports from an exporting to an importing country.

[7]I accessed this dataset through the Reuters Eikon terminal.

[8]Note that the dataset does not include freight rates for every possible importer-exporter pair. Freight rates for exporter-importer pairs not covered by Waterborne are imputed based on a regression model linking the freight rate to the distance between two ports.

able for the period 2004-2017 from Bloomberg.[9] I corroborate this with data with annual industry reports provided by the GIIGNL (The International Group of Liquefied Natural Gas Importers). In addition, for the vast majority of contracts, I have collected press releases and newspaper articles announcing the signing of the contract. This serves two purposes: it provides a way to independently verify the existence of the contract, and it allows me to construct a key variable: the date when the contract was signed, which is not available in the main source dataset (which only records the data when trade begins). The eventual dataset consists of every LNG contract exceeding 4 years that were signed in this industry, as well as the majority of shorter term contracts. Individual spot transactions (e.g. for a single cargo of LNG) are not included in this dataset.

For each contract, I observe the contract quantity, year during which contract was signed, and the contract duration (i.e. the contract start year and end year).[10] I observe the identity of the buyer, the identity of the seller, the import country and the export country. About 10% of the contracts (mostly in the latter part of the sample) are “portfolio" contracts where the seller is free to supply the LNG from anywhere in the world where it can gain access to LNG; in such cases the export country is recorded as “Unspecified". I observe whether the contract is signed on a free-on-board (FOB) basis (which loosely speaking means that the buyer is responsible for the shipping) or a delivered ex-ship (DES) basis (which means that the seller is responsible for the shipping). The dataset also records whether the contract is an extension of an earlier contract, or whether the contract is part of a “swap" arrangement with another contract.[11]

There are a number of contract details that I do not observe. The most important of these is the pricing formula, which indexes the price of LNG to the price of a selected energy benchmark (such as the oil price or the domestic price of natural gas). I also do not observe the “take-or-pay" share of the contract, whether the contract includes a destination restriction (which as described above restricts the ability of the buyer to divert the cargo to a third destination without permission from the seller), and any “diversion" clauses that specify how the gains from “diverting" LNG cargoes to a third party are to be divided between the buyer and the seller. Each of these details - contract pricing formula, the take-or-pay share, destination restrictions and diversion clauses - are not only missing from my dataset, but are also typically confidential and known only to the parties that are signatory to the contract. They are unlikely to be known with certainty by firms that are not part of the contract.

Table 1 contains summary statistics on key variables used in the analysis. Trade flows and

[9]The dataset was accessed through the Bloomberg terminal.

[10]Note that contracts are typically signed an average of two to three years prior to the period when trade begins. For an LNG contract beginning in 2010, for example, the buyer and seller would typically sign a sale and purchase agreement by 2007 or 2008.

[11]An example of a swap is when buyer A originally has a contract with seller A’, and buyer B originally has a contract with buyer B’. A swap deal would mean that buyer A now purchases from seller B’, while buyer B now purchases from seller A’. This kind of arrangement, though, is relatively rare in the industry.

Table 1: Summary Statistics

	Obs.	Mean	S.D.	Min	Max
Exporter-Importer-Year (ijt)					
Spot Trade (mt)	6,406	0.10	0.39	0	7.70
Contracted Trade (mt)	6,406	0.35	1.47	0	23.90
Re-exports (mt)	2,096	0.01	0.05	0	0.78
Shipping Cost (US$/MMBtu)	9,812	1.30	0.87	0.06	5.08
Exporter-Year (it)					
Export Capacity (mtpa)	238	14.10	15.21	0	77.00
Total Exports (mt)	238	12.31	14.55	0	79.63
Total Spot Exports (mt)	238	2.81	4.04	0	25.11
Importer-Year (jt)					
Total Imports (mt)	359	8.24	15.45	0	89.19
Total Spot Imports (mt)	359	1.94	3.27	0	25.74
Spot Prices (US$/MMBtu)	317	8.65	3.65	2.52	16.59

1. All trade variables (spot trade, contracted trade, total exports, total imports etc.) are measured in million tonnes or mt. Capacity is measured in million tonnes per annum, or mtpa.
2. Spot prices and shipping costs are measured in US$/MMBtu.

shipping costs are defined at the exporter-importer-year level; export capacity and total exports are defined at the exporter-year level and spot prices and total importers are defined at the importer-year level. A key fact to note is the prevalence of zeros in the trade data: out of 6,406 observations at the exporter-importer-year level, only 1,563 observations feature positive spot trade flows and 1,208 observations feature positive long-run contracted flows. Similarly, out of 2,096 observations at the re-exporter-importer-year, 275 observations feature positive re-exports.[12] The empirical analysis described latter deals with the censoring of the data induced by the presence of zeros.

3 Descriptive Evidence of Market Power

A key goal of the paper is to detect and quantify market power in the LNG spot industry. I now present some stylized facts and reduced-form evidence that, taken together, suggest sellers exercise market power in the LNG industry.

I first discuss a few features of the industry that make it possible in principle for LNG suppliers to exercise market power and price discriminate on the spot market. First of all, the LNG spot market is highly concentrated. Figure 4 plots the four-firm concentration ratio over time. It shows

[12] A re-exporter is defined as any LNG importer that has the capability of re-exporting LNG: not all LNG import terminals are equipped to engage in re-exports.

Figure 4: Industry concentration (four-firm concentration ratio) over time

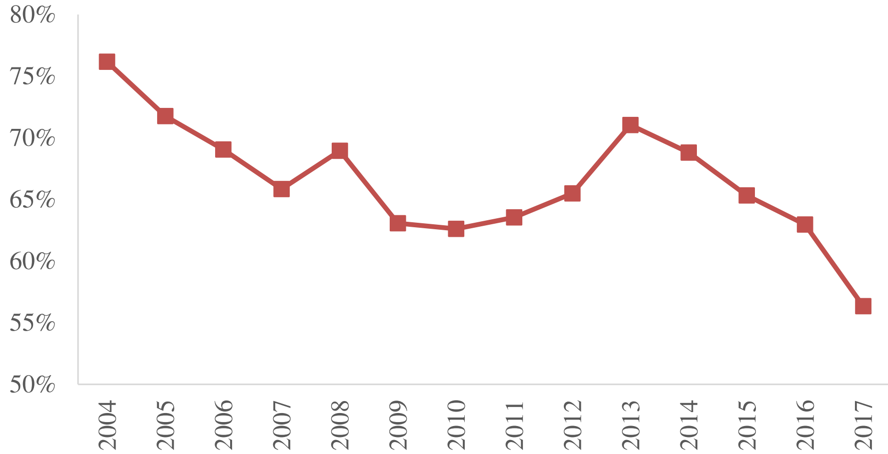

Note: The figure plots the four-firm concentration ratio (CR4) for the LNG spot market over time, calculated as the share of spot and short-term trade accounted for by the top four exporting countries. The CR4 is calculated using spot trade data from GIIGNL.

Figure 5: Shipping costs ($/MMBtu) for major LNG routes

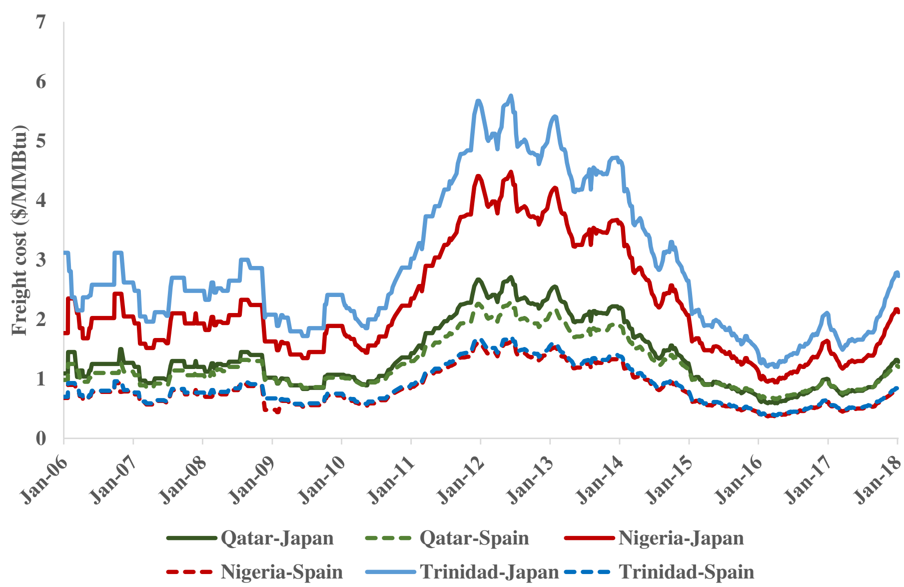

Note: The figure plots the shipping cost for major LNG trading routes. Shipping costs are denoted in $/MMBtu. Data from Waterborne.

that the top four exporters account for over 60% of spot trade in most years. decreasing only from 2014 onward, driven by entry of new LNG suppliers in Australia and the US. The high degree of concentration facilitates the ability of sellers to exercise market power in the spot market.

Secondly, shipping costs are important in this industry. High shipping costs cause markets to become segmented from one another and make it feasible in theory for sellers to price discriminate across markets. As Figure 5 illustrates, the freight rate varies considerably over time, driven by both changes in the crude oil price (which determines the fuel cost incurred in LNG shipping) and changes in LNG demand: during periods of high LNG demand, demand for LNG shipping increases and this leads to higher freight rates. When shipping costs are high, sellers face less competition both from faraway sellers and faraway buyers who can engage in arbitrage behavior. For example, as Figure 5 shows, shipping costs were at historically high levels between 2012 and 2014. During this period, Qatar (the largest LNG producer) faced less competition when selling to Asia (the high-demand market during this period), as some of its major competitors (Nigeria and Trinidad) incurred much higher transportation costs when selling into Asia. Moreover, the high transportation costs increased arbitrage costs for European buyers looking to re-export LNG from Europe to Asia. This made it more feasible for Qatar as well as other sellers located closer to Asia to charge higher prices when selling to Asian buyers.

Third, several institutional characteristics of the LNG industry impede arbitrage and thereby facilitate price discrimination by sellers.[13] A large proportion of LNG is traded under long-term contracts, and the buyers of contracted LNG could potentially engage in arbitrage. However, the existence of destination restrictions in long-term contracts impedes arbitrage, since the buyer has to take physical delivery of the LNG at their terminal and does not have the option of diverting the LNG cargo. Contracts can sometimes have clauses to deal with scenarios where the seller and the buyer wish to divert cargoes to a third destination, but diversions in practice are impeded by the difficulty of agreeing terms on profit sharing (GIIGNL, 2012). Moreover, a seller who is already selling to the third destination has little incentive to allow such cargo diversions, since the result would be that the seller faces competition in that market from their own contracted buyer.

LNG buyers do have the ability to receive a contracted LNG cargo at their liquefaction terminal, reload the LNG onto a new tanker and then re-sell the cargo to a new destination. The costs of engaging in such re-exports can, however, be high. They include not just the shipping cost between the original and the new destination, but also a fee paid to the terminal operator as well as a host of logistical constraints that further add to the costs of re-exporting.[14] As a result, an industry report on re-exports argues that “These logistical factors mean that the true cost of reload is well above

[13]See Ritz (2014) for a further discussion of barriers to LNG arbitrage.

[14]A would-be reloader needs to find both a cargo that is available for re-exports and a buyer who is willing to take the cargo within a few days. Storing LNG is costly and due to energy lost via boil-off it is typically not feasible to store LNG for long periods of time.

the direct cost paid to the terminal operator. In other words a premium well in excess of shipping costs is required to incentivize the re-exporting of gas."[15]

The above discussion suggests that the *possibility* for sellers to exercise market power exists in this industry. At the same time, it is not a priori obvious just from these facts that sellers would exercise market power. The firms involved in LNG exporting typically have close ties to national governments: indeed, in the majority of LNG exporting countries, LNG exporting companies are majority-owned by national oil and gas companies.[16] Thus it is not immediately apparent that LNG exporters' behavior on the spot market would conform to short-run profit maximization as in standard oligopolistic models. One could imagine too that sellers might have motives to refrain from price discriminating as a way to build relations with buyers, given that most of their sales still come from long-term contracts. Finally, though the LNG industry is significantly less concentrated on the buyer side than on the seller side, there do exist large buyers, unlike in standard settings of oligopolistic behavior. As such, whether firms exercise market power and price discriminate is ultimately an empirical question that needs to be answered by looking at how firms actually behave on the spot market, rather than assumed from the outset.

I now describe the recent evolution of spot prices and spot trade flows in the industry, and argue that these patterns are difficult to reconcile with competitive behavior on the part of the sellers. Figure 6 shows spot prices in various LNG importing regions. The feature that stands out from this figure is the presence of significant and systematic spot price differentials across regions, especially during periods when the LNG market is tight. A particularly striking illustration of this is the period between mid-2011 and end-2013, when there was a large spike in Japan's LNG demand following the Fukushima nuclear disaster. Asian spot prices, as well as spot prices in Latin America, remained an average of $5/MMBtu higher than European prices during this period. The prices converge again only in late 2014 and early 2015, driven by a combination of reduced LNG demand in Asia, lower transportation costs (driven by the oil price collapse) and addition of new export capacity starting from 2015. In addition to the post-Fukushima period from 2011 to 2013, there have been other periods when prices in different regions sharply diverged from one another. For example, between January 2007 and July 2008, another period of tight demand, Asian spot prices were about $3.5/MMBtu higher on average than European spot prices.

In a competitive spot market with no frictions, faced with these divergent prices, sellers would sell only to the destination with the highest price net of transportation costs. This would suggest that during these periods of high Asia-Europe spot price differentials, spot exports should be mostly directed to Asia, as even after accounting for transportation costs most sellers received higher prices from selling to Asia than to Europe. As Figure 7 shows, though, during the "boom" period of 2011-

[15]See http://www.timera-energy.com/will-european-lng-reloads-continue/.

[16]For example, Qatar Petroleum has a 70% stake in Qatargas and Rasgas, the two LNG producing companies in Qatar.

Figure 6: LNG Spot Prices in Different Regions

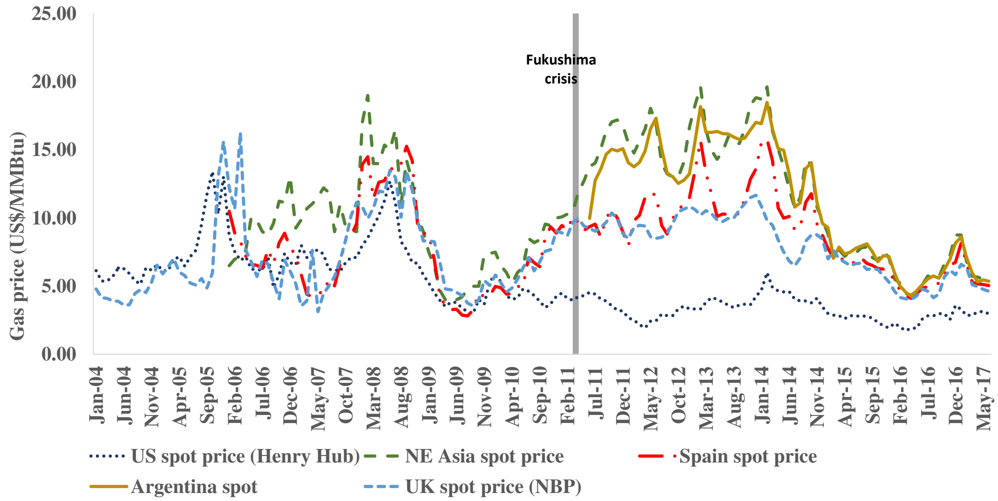

Note: The figure plots monthly spot LNG and gas prices in different regions. The US price is the price of natural gas traded on the Henry Hub. The UK price is the price of natural gas traded on the NBP Virtual Trading Point. The price in North-east Asia is a benchmark spot price for the region (comprising Japan, Korea, China and Taiwan) reported by Reuters. Finally the spot prices in Spain and Argentina are reported by Waterborne LNG.

Figure 7: Short-run and spot LNG imports to different regions

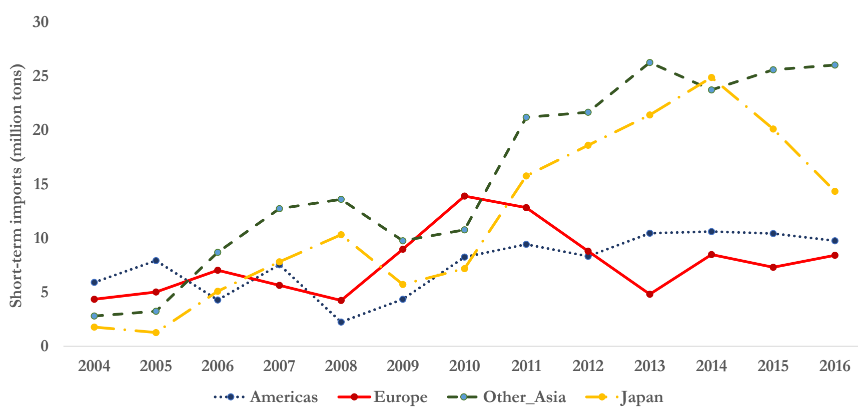

Note: Short-run and spot LNG imports refers to LNG imports obtained from spot transactions or from contracts that are less than or equal to four years in length. The source data is from GIIGNL.

14 Europe continued to import a significant amount of LNG on the spot market.[17] Moreover, much of Europe's spot market imports during this period came from sellers that would have received higher prices from selling to Asia. For example, in 2012, 46% of Europe's spot market purchases were from Qatar, and a further 12% were from Egypt and Algeria. This is in spite of the fact that these countries faced roughly similar shipping costs of selling LNG to Asia and Europe (see Figure 5 for shipping costs from Qatar to Asia and Europe), meaning the marginal cargo garnered a premium of $2-4/MMBtu if it were sent to Asia rather than Europe. The fact that sellers in Middle East and North Africa continued to sell to Europe is inconsistent with competitive behavior on the part of LNG exporters.

The pattern of sellers selling to multiple markets at different received prices suggests that sellers deviate from perfectly competitive behavior. However this still does not establish whether these deviations are the result of market power or some other sources, such as unobserved trade frictions. I now present a reduced-form test for whether sellers exercise market power. The test is designed to distinguish the role of market power from other trade frictions. The intuition behind the test is that if sellers exercise market power, small and large sellers should allocate their sales differently across markets. LNG sellers have the choice of selling to several markets, each of which has different demand shocks. If the seller is small in size, it should behave almost competitively and sell most of its output to the market from which it receives the highest price (net of shipping costs). If the seller is large and exercises market power, though, it has an incentive to withhold LNG from the market with high demand and sell more LNG to the market with low demand, creating a wedge in prices between the two destinations. Thus, if sellers exercise market power, we should see large sellers on the spot market respond less strongly to price differentials than small sellers. Under the alternative hypothesis of perfectly competitive behavior, even with trade frictions, both small and large sellers should be equally responsive to spot price differentials.

To formalize this intuition, we study how the relationship between the share of a seller's spot LNG output that it allocates to destination j and the price that the seller receives from destination j depends on a seller's size on the spot market (as measured by its total spot sales). Let i denote sellers, j denote buyers and S_{ij} denote the spot sales from country i to country j. Let $S_i = \sum_j S_{ij}$ denote country i's total spot sales. Finally let p_{ij} denote the spot price that country i receives from country j net of the shipping cost between i and j.[18]. We carry out the following regression:

$$ln(S_{ijt}/S_{it}) = \alpha + \beta ln(p_{ijt}) + \gamma ln(p_{ijt}) * ln(S_{it}) + X_{ijt} * \theta + \eta_{ijt} \quad (1)$$

[17] Even more significantly, Europe continued to import large amounts of contracted LNG, suggesting that long-term contracts also impede the market's response to demand shocks. I explore the rigidity of contracted trade in responding to demand shocks further in the ongoing research project mentioned in the introduction.

[18] If the price at country j equals p_j and the shipping cost between i and j, then $p_{ij} = p_j - d_{ij}$.

In this regression, β captures how responsive spot sales are to price differentials: the larger β is, the greater the proportion of a seller's spot output that it allocates to market j. The magnitude of β will generally be larger the lower the magnitude of unobserved trade frictions. The key coefficient of interest is γ, which measures how the responsiveness of spot sales to price differentials differs by seller size. $\gamma = 0$ would be consistent with the absence of seller market power, while $\gamma < 0$ implies that sellers with larger sales to a market respond less strongly to price differentials (since they have an incentive to withhold LNG from a market to which they already sell a lot).

Table 2: Reduced Form Evidence for Market Power: Regression of Share of Spot Sales to market j on Price at market j

	(1) OLS	(2) OLS	(3) OLS	(4) OLS	(5) IV	(6) IV
	Dependent variable: log(share of spot sales to market j)					
β*log(p_{ijt})	0.44***	0.72***	0.99***	0.84***	0.60***	0.62***
	(0.070)	(0.063)	(0.096)	(0.10)	(0.19)	(0.19)
γ*log(p_{ijt})*log(S_{it})		-0.28***	-0.27***	-0.25***	-0.24***	-0.26***
		(0.014)	(0.014)	(0.022)	(0.027)	(0.040)
Year fixed effects			Y	Y	Y	Y
Exporter fixed effects				Y	Y	Y
N	1457	1457	1457	1457	900	898
R^2	0.026	0.25	0.26	0.31	0.33	0.32
1st-stage F-statistic					116	109

Note: Each observation is an exporter-importer-year pair. Total LNG spot exports from country i to country j in year t are measured in million tonnes. The price i receives from j is measured in \$/MMBtu. The variable size$_i$ refers to seller i's total spot sales. In column (5), we instrument for log(p_{ijt}) with log(electricity consumption from fossil fuels$_{jt}$) and log(coal price$_{jt}$). In column (6), we instrument for log(p_{ijt}) with log(electricity consumption from fossil fuels$_{jt}$) and log(coal price$_{jt}$) and instrument for log(S_{it}) with log(Spot capacity$_{it}$), where seller i's spot capacity is defined as its total export capacity minus its contracted exports. The 1st-stage F-statistic in columns (5) and (6) is the Cragg-Donald Wald F statistic. Statistical significance at the 10%, 5%, and 1% levels are denoted with *, **, and ***, respectively.

Table 2 shows the estimated coefficients from this regression. As expected, the coefficient on the logarithm of prices is positive: holding all other prices fixed, if a demand shock causes the price at destination j increase by 10%, the share of spot output sold to that market increases by 4.4%. However the coefficient on the interaction term between price and size is negative, implying that the larger the total spot sales of seller i, the lower is seller i's quantity response to a price increase in destination j. Moreover, comparing columns (1) and (2), there is a noticeable increase in the R^2

when we include this term, suggesting that a large part of the variation in spot sales is driven by market power.

The finding survives the inclusion of year fixed effects that control for year-to-year variation in aggregate demand conditions and exporter fixed effects (columns (3) and (4)). One might worry about possible endogeneity of both spot prices and the firm's size (as measured by its total sales). In column (5), we instrument for the price p_{ijt} with demand shifters at destination j: electricity consumption from fossil fuels and the regional price of coal (which is a competing fuel). In column (6) we additionally instrument for the seller's total spot sales S_{it} with its spot capacity in period t (i.e. its capacity left over after fulfillment of contracts). The idea is that although a seller's total spot sales may adjust in response to demand conditions in market j, the seller's spot capacity is determined by investment and contracting decisions made many years in the past, and hence plausibly exogenous to current demand shocks. The estimated γ is very similar when we instrument for price and seller size. The finding that $\gamma < 0$ is difficult to reconcile with competitive behavior but is consistent with the hypothesis of market power.

Finally, some additional evidence suggestive of seller market power is provided by Figure 8, which plots the average number of buyers each seller sells to on the spot market against its total spot market sales. There is a positive correlation between seller size and the number of buyers to which the seller sells, suggesting that large sellers are more likely to price discriminate by splitting their sales across multiple buyers, while small sellers generally sell to a small number of buyers that offer the highest price.

Figure 8: Relationship between total spot sales and number of countries that a seller exports to

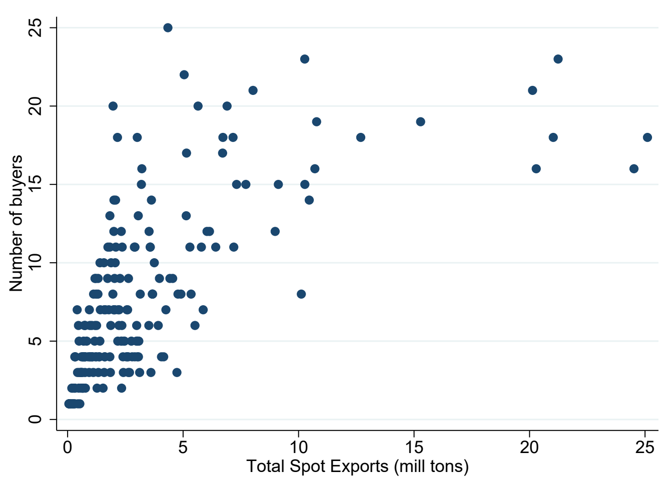

Note: The x-axis plots total spot sales of each seller. The y-axis plots the number of buyers that each seller sells to on the spot market

4 Model

In this section, I describe a static model of trade and pricing in the liquefied natural gas industry. The model features a set of sellers and buyers that are located in geographically distinct locations. Capacity-constrained sellers produce a homogeneous good that they can sell to the buyers, but only by incurring a shipping cost that is increasing in the geographical distance between the seller and the buyer. Each buyer has a downward-sloping curve with a stochastic intercept that reflects buyer-specific idiosyncratic demand shocks.

Trade takes place via both contracts and a spot market. Contracted quantities are fixed in the short-run and cannot adjust in response to the demand shocks. In the spot market, by contrast, sellers observe demand shocks prior to making any decisions. Under the hypothesis of Cournot competition, sellers choose their own spot quantities taking as given the spot quantities chosen by their rivals as well as all contracted quantities. Under the hypothesis of perfect competition, sellers choose their spot quantities taking equilibrium prices as given. The model nests both of these models of seller behavior.

Capacity constraints play a key role in the model. In the absence of capacity constraints, a seller's decision in one market can be analyzed independently of what the seller chooses to do in other markets. Because of capacity constraints, though, a seller's choice of spot sales in one market will necessarily depend on the seller's chosen spot sales in other markets. This causes markets to be inter-connected, with demand shocks in one market having ripple effects in other markets as well.

4.1 Model Description

Buyers

There are J buyers indexed by $j = 1, ..., J$, that are geographically differentiated from one another. The demand curve for buyer j is given by the following equation:

$$Q_j = Q_j^d(p_j, \varepsilon_j) \tag{2}$$

where ε_j denotes demand shocks. The inverse demand curve is $p_j = P_j(\varepsilon_j, Q_j)$. Note that Q_j denotes buyer j's *total* purchases of LNG, including both contracted LNG and spot LNG.

In the empirical analysis I will specialize to the case of linear demand:

$$Q_j^d(p_j, \varepsilon_j) = \varepsilon_j - bp_j \tag{3}$$

With linear demand, the inverse demand curve is: $p_j = \frac{1}{b}(\varepsilon_j - Q_j) = \beta(\varepsilon_j - Q_j)$ where β is defined as $\frac{1}{b}$.

Sellers

There are N sellers indexed by $i = 1, ..., N$. Each seller has capacity k_i that determines its marginal costs of production.

The quantity of LNG sold by seller i to buyer j equals $q_{ij} = S_{ij} + q^c_{ij}$ where S_{ij} denotes spot sales and q^c_{ij} denotes contracted sales. Thus the total amount of LNG produced by seller i equals $q_i = \sum_j q_{ij} = \sum_j (S_{ij} + q^c_{ij})$, while the total quantity imported by buyer j equals $Q_j = \sum_i q_{ij} = \sum_i (S_{ij} + q^c_{ij})$.

A key assumption I make is that contracted quantities q^c_{ij} are fixed in the short-run and cannot adjust in response to the demand shocks. This reflects the empirical reality in the LNG industry that most contracts are signed well in advance of when trade takes place. In contrast, S_{ij} can be chosen by the sellers every year after observing the demand shocks, and is the key strategic variable in the model of the spot market.

It should be noted that when taking the model to data, I use the actual contracted volume traded between exporter i and the importer j, rather than the contracted volume that they had originally agreed to. Thus my approach does allow for some flexibility to the contracted trade volume. What I do not allow for is the possibility that the contracted trade volume adjusts in response to the spot trade volumes chosen by the exporters, since in the model all agents take q^c_{ij} as given when making any decisions.

Costs of production and sales

Firms incur costs in producing LNG, and then in selling the LNG to the buyers.

Let $C(q_i, k_i)$ denote the cost of production, which I will assume is continuously differentiable in both arguments. The inclusion of k_i in the cost function reflects the fact that capacity constraints are important in this industry (as evidence by Figure 3) and sellers with larger capacity are able to export more LNG.

In addition, the firm incurs costs of trading that differ by buyer, which consists of three components. Each unit of LNG costs d_{ij} to ship from i to j, where the shipping cost d_{ij} is increasing in the distance between seller i and buyer j and fluctuates by year.[19] I allow the costs of selling LNG to buyer j to be shifted by observable characteristics x_{ij}. For example, there may be costs of building a relationship which are only incurred when a seller and buyer first engage in trade; in this case, x_{ij} would include an indicator for whether or not i and j have traded in the past. Finally the per-unit cost of trade is shifted by an idiosyncratic match cost ν_{ij} that is drawn i.i.d. from the $N(0, \sigma^2)$ distribution. The match cost reflects idiosyncratic trade frictions that might change the

[19] d_{ij} is directly observed in the dataset,. In practice d_{ij} is a linear function of the nautical distance between seller i and buyer j, since the cost of shipping increases linearly with the journey time. The slope of this linear function as well as the intercept term differ by year, reflecting prevailing LNG shipping rates at the time.

trade cost between seller i and buyer j. For example, if scheduling constraints are important, and seller i happens to be the only seller that can deliver a cargo to buyer j at the time when buyer j requires the cargo, then this will be reflected in a lower value of ν_{ij} for that buyer-seller pair.

Thus firm i's cost function is:

$$C_i(q_{i1},..q_{ij},...q_{iJ},k_i) = \underbrace{C(q_i,k_i)}_{\text{Cost of production}} + \underbrace{\sum_j (d_{ij} + \phi(x_{ij}) + \nu_{ij})q_{ij}}_{\text{Cost of trading}} \tag{4}$$

Re-sales

Buyers can engage in arbitrage ("re-sales"): they can re-sell any imported quantities to other buyers. In order to engage in re-sales, they have to pay both shipping costs and an additional resale cost κ per unit. This additional cost κ includes a direct cost levied by terminal operators as well as a number of logistical costs associated with arranging a resale to occur within a few days after the cargo is received. I assume resale costs are convex in quantity so that the marginal cost increases in the quantity re-sold. Thus the cost of selling $q_{j'j}$ units of LNG from buyer j' to buyer j is $C^{resale}(q_{j'j}) = (p_{j'} + d_{j'j} + \kappa)q_{j'j} + \delta^r q_{j'j}^2$.

When resales are zero, the marginal cost of re-selling gas from buyer j' to buyer j is given by $p_{j'} + d_{j'j} + \kappa$: this equals the cost of purchasing the LNG on the spot market added to the shipping cost and the additional re-sale cost. As re-sales increases the marginal cost increases.

Buyer j' will only re-sell to buyer j if p_j exceeds $p_{j'} + d_{j'j} + \kappa$: so resales will only take place when price differentials exceed the sum of shipping cost and resale cost. I assume that sellers cannot buy and then re-sell LNG: only regasification terminals (which are the buyers' terminals) have the ability to take LNG, store it for a few days and reload.

Information structure

I assume that the match cost between seller i and buyer j is seller i's private information that seller i observes *only* when choosing how much to sell on the spot market to buyer j. Thus, at the time of choosing S_{ij}, seller i does not observe the match costs of its rivals, nor does it observe its own match costs for choosing spot sales to other buyers $k \neq j$.

The assumption of privately observed match costs is meant to capture the fact that there may be factors affecting the costs of trading between seller i and buyer j that other sellers do not observe. Moreover, in practice sales to different buyers take place at various points in time and when selling to buyer j, seller i might not yet know the full cost of selling to another buyer k. This is the motivation behind the assumption that seller i only sees ν_{ij} when choosing spot sales to buyer j, but does not observe ν_{ik} for $k \neq j$.

It is also possible instead to assume that the environment is characterized by common knowledge: the full vector of ν_{ij} is observed by *every* seller when choosing spot sales. The qualitative predictions of the model are very similar if I make this alternative assumption; however estimating the model with common knowledge is challenging because computing the likelihood requires numerically solving for high-dimensional integrals, which is analytically infeasible when there are many sellers and buyers. The assumption of private information that I adopt instead aids in analytical tractability.

Other than the match cost, every other relevant variable is common knowledge and observed by all sellers at the time when they choose spot sales. This includes the set of buyers and sellers; contracts, q_{ij}^c; capacity k_i; shipping costs d_{ij}; trade cost shifters x_{ij}; and demand shocks ε_j.

Seller competition

As discussed earlier, contracted sales are treated as predetermined and cannot be changed. Taking contracted sales and capacity as given, how will spot prices and spot trade flows be determined?

I consider two possible models of seller behavior. The first model allows sellers to behave strategically in the spot market. The market is characterized by a Bayes-Cournot equilibrium, where sellers compete in quantities. Each seller's strategy is to choose $S_{ij}(\nu_{ij})$, taking as given the strategies chosen by its rivals $\{S_{-i,k}(\nu_{-ik})\}_{k=1}^{J}$ as well as its own strategies in other markets $\{S_{ik}(\nu_{ik})\}_{k\neq j}$. A second model of seller behavior is perfect competition. Under this model, sellers choose quantities $S_{ij}(\nu_{ij})$ taking expected prices as given, without accounting for their price impact. When estimating the model, I do not impose either of these assumptions, and instead formally test between the two models.

Throughout I assume that buyers are price-takers, both when they purchase LNG on the spot market, and when they re-sell LNG. Oligopsony is unlikely to be a first-order concern, because there are many more importing countries than exporting countries, and the buyers' side of the market tends to be much more fragmented: there are usually several importing firms in each importing country. In addition, resales are relatively small relative to overall spot sales so it is unlikely a buyer who is re-selling LNG will deviate very much from perfectly competitive pricing.

4.2 Equilibrium

Market clearing

Each of the J markets clears separately and simultaneously. The market clearing price vector $p^* = (p_1, ..., p_J)$ is characterized by the following set of equations:

$$Q^d(p_j^*, \varepsilon_j) = \sum_{i=1}^{N} q_{ij}^c + \sum_{i=1}^{N} S_{ij}(\nu_{ij}) \;, \forall j \tag{5}$$

Optimal behavior

I begin by assuming that sellers choose spot market quantities strategically and take into account the effect their decisions have on the market price: in other words, the equilibrium is characterized by Cournot competition. I then show how to nest perfect competition within the model.

In a Bayes-Cournot equilibrium, each seller i, after observing ν_{ij}, takes as given rival optimal strategies $\{S_{-i,k}(\nu_{-ik})\}_{k=1}^{J}$ as well as its own strategies in other markets $\{S_{ik}(\nu_{ik})\}_{k\neq j}$, and chooses its own spot sales to market j, $S_{ij}(\nu_{ij})$, in order to maximize its expected profits across all markets:

$$max_{\hat{S}_{ij}} E\left[\sum_{k=1}^{J} p_{ik}^c q_{ik}^c + \sum_{k\neq j}^{J} p_k^* S_{ik}(\nu_{ik}) + p_j^* \hat{S}_{ij}(\nu_{ij}) - C(q_i, k_i) - \sum_k (d_{ik} + \phi(x_{ik}) + \nu_{ik}) q_{ik}\right] \tag{6}$$

where recall that $q_i = \sum_k (q_{ik}^c + S_{ik})$

The revenue from contracts $\sum_{k=1}^{J} p_{ik}^c q_{ik}^c$, the spot revenue from selling to other buyers $\sum_{k\neq j}^{J} p_k^* S_{ik}(\nu_{ik})$ and the shipping/match costs from sales to other buyers are unaffected by the choice of S_{ij} and hence the above optimization problem can equivalently be expressed as:

$$max_{\hat{S}_{ij}} E\left[p_j^* \hat{S}_{ij}(\nu_{ij}) - C(q_i, k_i) - (d_{ij} + \phi(x_{ij}) + \nu_{ij}) \hat{S}_{ij}(\nu_{ij})\right]$$

$$max_{\hat{S}_{ij}} E(p_j^* | \hat{S}_{ij}(\nu_{ij})) \hat{S}_{ij}(\nu_{ij}) - E(C(q_i, k_i) | \hat{S}_{ij}(\nu_{ij})) - (d_{ij} + \phi(x_{ij}) + \nu_{ij}) \hat{S}_{ij}(\nu_{ij}) \tag{7}$$

where equation (7) makes it clear that the seller faces uncertainty both over the price in market j as well as its total costs (since it does not know exactly how much it will sell in other markets).

The first-order condition satisfied by the optimal quantity $S_{ij}(\nu_{ij})$ is:

$$E(p_j^* | S_{ij}(\nu_{ij})) + S_{ij}(\nu_{ij}) \frac{\partial E(p_j^* | S_{ij}(\nu_{ij}))}{\partial S_{ij}} - \frac{\partial E(C(q_i, k_i) | S_{ij}(\nu_{ij}))}{\partial S_{ij}} - (d_{ij} + \phi(x_{ij}) + \nu_{ij}) \leq 0 \tag{8}$$

with equality iff $S_{ij}(\nu_{ij}) > 0$.

Notice that because of capacity constraints, the seller's choice of S_{ij} depends on how much he expects to sell in other markets. If the seller expects to sell high quantities in some other market (due to say a demand shock in that market), then $\frac{\partial E(C(q_i,k_i)|S_{ij}(\nu_{ij}))}{\partial S_{ij}(\nu_{ij})}$ will be higher and so the seller will tend to sell less in market j. Thus although the firm does not perfectly coordinate its sales across other markets (due to imperfect information), it takes into account cross-market externalities when choosing sales. In particular, as alluded to before, demand shocks in one market will affect firms' sales decisions across all markets.

If firms were instead to behave non-strategically and act as price-takers, they do not take into account the effect of their choice of S_{ij} on the equilibrium price . The first-order condition would then change to:

$$E(p_j^*|S_{ij}(\nu_{ij})) - \frac{\partial E(C(q_i,k_i)|S_{ij}(\nu_{ij}))}{\partial S_{ij}} - (d_{ij} + \phi(x_{ij}) + \nu_{ij}) \leq 0 \tag{9}$$

with equality iff $S_{ij}(\nu_{ij}) > 0$.

Both cases can be compactly summarized by introducing a conduct parameter $\theta_i^c \in \{0,1\}$ (Bresnahan, 1982), as in equation (10). When $\theta_i^c = 0$, firms behave competitively. When $\theta_i^c = 1$, firms engage in Cournot competition:

$$E(p_j^*|S_{ij}(\nu_{ij})) + \theta_i^c S_{ij}(\nu_{ij}) \frac{\partial E(p_j^*|S_{ij}(\nu_{ij}))}{\partial S_{ij}} - \frac{\partial E(C(q_i,k_i)|S_{ij}(\nu_{ij}))}{\partial S_{ij}} - (d_{ij} + \phi(x_{ij}) + \nu_{ij}) \leq 0 \tag{10}$$

A key goal of the estimation exercise will be to estimate equation (10) to see if the actual behavior of sellers is better characterized by perfect competition or by Cournot competition.

Equilibrium re-sales

Since buyers are price-takers when reselling, equilibrium re-sales are characterized by the following first-order condition:

$$\forall j, j', q^s_{j'j} = max\big(0, \frac{1}{\delta^r}(p_j - p_{j'} - d_{j'j} - \kappa)\big) \text{ if } \sum_j q^s_{j'j} < Q_{j'} \tag{11}$$

The possibility of resales means that the residual demand curve facing any of the producers selling to a buyer j becomes more elastic once the price difference between any two regions exceeds $d_{j'j} + \kappa$ and resales to buyer j become possible. The buyer of course cannot resell more than the amount of LNG he has purchased which is why we impose the condition $\sum_j q^s_{j'j} < Q_{j'}$.

In Appendix A, I describe Monte Carlo simulations of a simplified version of the above model that I carried out in order to illustrate the effect of market power in this model. In the simplified version of the model, trade costs are deterministic and marginal costs are linear until production hits a binding capacity constraint. I first consider a scenario where firms have large amounts of excess capacity. I show that market power results in both lower overall production and greater splitting of sales across buyers, with sellers more likely to sell to faraway buyers than under competitive behavior. I then consider a scenario where firms are highly capacity-constrained, so that they cannot adjust their total production. I show that in this case, market power still leads to lower welfare because the allocation of quantities across markets is inefficient. Sellers sell more of their output to faraway destinations under Cournot competition. When there is a demand shock in one market, sellers do not increase their sales to that market as sharply as under competitive behavior, which results in higher prices in the high-demand market than in the low-demand market and misallocation across the two markets.

5 Estimation

I now describe how to estimate the model described in Section 4. I allow the data to identify whether sellers behave strategically as in Cournot competition or non-strategically as in perfect competition. I do so by estimating a conduct parameter that nests various models of seller behavior, together with parameters describing seller costs. A precondition for estimating cost and conduct parameters is knowledge of the industry demand curve. As such I begin by describing how I estimate the demand curve, and then follow up by describing the method for estimating costs together with the conduct parameter. For now I abstract away from estimating buyer re-sale costs, but plan to incorporate resale cost estimation into future versions of the paper.

5.1 Estimation of Demand Curve

I estimate the following demand curve for importing country j:

$$Q_{jt} = Q_j^d(p_{jt}, X_{jt}, \varepsilon_{jt}) = \alpha - bp_{jt} + X_{jt}\gamma + \varepsilon_{jt} \tag{12}$$

Here X_{jt} denotes demand shifters. In the empirical analysis, this includes electricity consumption as well as country fixed effects. Liquefied natural gas is almost always converted back into gaseous form before it is consumed, and natural gas is primarily used for power generation and heating. Electricity consumption thus captures changes in electricity demand that translate into a shift in the demand curve for LNG. I use country fixed effects to capture differences in heating demand across countries as well as other time-invariant country-specific factors that affect LNG demand (for example the amount of piped natural gas the country has access to).

Estimation of equation (12) via ordinary-least squares runs into standard endogeneity problems because the spot price p_{jt} is positively correlated with the unobserved demand shock ε_{jt}: the greater

Table 3: Demand Curve Estimates

	(1) OLS	(2) IV	(3) IV	(4) IV
Spot Price	0.70***	-1.40*	0.091	-0.35***
	(0.24)	(0.80)	(0.58)	(0.13)
Electricity Consumption			0.0039***	0.016***
			(0.00099)	(0.0011)
Importer Fixed Effects	No	No	No	Yes
N	317	317	297	297
R^2	0.026	-0.21	0.057	0.96
1st-stage F statistic		19.60	33.74	34.87

Note: Each observation is an importer-year pair. The dependent variable is total LNG imports in country j in year t, measured in million tonnes. The spot price is measured in \$/MMBtu. In (2)-(4), I instrument for prices using total liquefaction capacity in period t and total electricity demand in period t excluding country j's own electricity demand.

country j's demand for LNG, the higher the equilibrium spot price in country j. I consider two instruments for p_{jt}. The first instrument, which varies only over time but not across j, is total liquefaction capacity (i.e. total export capacity) in the world in period t. The greater LNG capacity in period t, the higher is the supply of LNG and therefore the lower the price in period t. The identification assumption is that LNG export capacity is uncorrelated with idiosyncratic demand shocks today, after controlling for electricity consumption. The logic behind the instrument is that LNG terminals take many years to build, and at the time the decision to invest is made, it is difficult to foresee idiosyncratic demand shocks that are realized several years later. The modern history of the LNG industry is replete with instances where sellers make investments without fully anticipating how demand would evolve in the importing countries. For example, Qatar's massive capacity expansion in the 2000s was driven to a large degree by the expectation that the US would be a major importer of natural gas. However by the time all of Qatar's terminals came online, US demand for LNG had shrunk dramatically due to the shale gas boom, and instead Qatar ended up turning to Asian countries as its major buyers of LNG.

The second instrument, which varies over both t and j, is electricity demand in other countries. The argument for relevance of the instrument is that if other countries experience higher electricity demand, then overall demand for LNG rises, and the price paid by country j will increase, holding fixed country j's own demand shifters. The instrument is plausibly exogenous provided that country

j's idiosyncratic demand shocks, after controlling for its own electricity demand, are uncorrelated with electricity demand in other countries.

The demand curve estimates are shown in Table 3. Column (1) shows the OLS regression of demand on spot prices, without any controls. As expected, the coefficient is positive, reflecting the fact that price changes are mostly driven by shifts in demand rather than shifts in supply. In Column (2)-(4), I instrument for spot prices using total liquefaction capacity and total electricity demand in countries other than country j. The 1st-stage Cragg-Donald Wald F-statistic ranges from 19 to 35, suggesting that weak instruments are unlikely to be a concern. In the most flexible specification (4) where I control for both electricity demand and importer fixed effects, the coefficient on the spot price is negative and statistically significant, and implies a demand elasticity of around 0.33 for the average observation. The coefficient on electricity consumption is also positive and statistically different from zero, and implies that the elasticity of LNG demand with respect to electricity consumption is around 0.82 for the average observation.

5.2 Estimation of Cost and Conduct Parameters

The starting point for the estimation of the cost and conduct parameters is the first-order condition (10), which I have repeated here:

$$E(p_j^*|S_{ij}(\nu_{ij})) + \theta_i^c S_{ij}(\nu_{ij}) \frac{\partial E(p_j^*|S_{ij}(\nu_{ij}))}{\partial S_{ij}} - \frac{\partial E(C(q_i,k_i)|S_{ij}(\nu_{ij}))}{\partial S_{ij}} - (d_{ij} + \phi(x_{ij}) + \nu_{ij}) \leq 0$$

To operationalize this for estimation, I need to specify a functional form for production costs. I assume that the production cost is quadratic, with the slope of marginal costs allowed to depend on k_i. The idea behind this is that firms with small k_i will become capacity-constrained very quickly, meaning their marginal costs increase rapidly with q_i. By contrast, firms with large k_i experience a more gradual increase in marginal costs as they raise q_i.

$$C(q_i,k_i) = c_i q_i + \frac{\delta(k_i)}{2} q_i^2 \tag{13}$$

I show in Appendix B that under the assumption of linear demand and quadratic costs, the first-order condition can be written as:

$$\beta\big(\varepsilon_j - Q_j^c - \sum_{i' \neq i} E(S_{i'j}(\nu_{i'j}))\big) - \big(c_i + \delta(k_i)\big(q_i^c + \sum_{k \neq j} E(S_{ik}(\nu_{ik}))\big) + d_{ij} + \phi(x_{ij})\big) - \\ \big(\beta(1+\theta_i^c) + \delta(k_i)\big) S_{ij}(\nu_{ij}) - \nu_{ij} \leq 0 \tag{14}$$

Let A_{ij} be defined as follows:

$$A_{ij} = \beta\big(\varepsilon_j - Q_j^c - \sum_{i' \neq i} E(S_{i'j}(\nu_{i'j}))\big) - \big(c_i + \delta(k_i)\big(q_i^c + \sum_{k \neq j} E(S_{ik}(\nu_{ik}))\big) + d_{ij} + \phi(x_{ij})\big) \quad (15)$$

Intuitively we can think of A_{ij} as representing a measure of the profitability of selling to buyer j: the greater A_{ij} is, the more likely seller i is to sell to buyer j.

The FOC can be written as:

$$A_{ij} - \big(\beta(1+\theta_i^c) + \delta(k_i)\big)S_{ij}(\nu_{ij}) - \nu_{ij} \leq 0 \quad (16)$$

Define

$$S_{ij}^* = \frac{1}{\beta(1+\theta_i^c) + \delta(k_i)}(A_{ij} - \nu_{ij}) \quad (17)$$

Then the firm's policy function is:

$$S_{ij} = max(S_{ij}^*, 0) \quad (18)$$

In particular the firm chooses $S_{ij} > 0$ if and only if $\nu_{ij} < A_{ij}$.

The firm's expected sales, before observing ν_{ij}, are given by:

$$\begin{aligned} E(S_{ij}(\nu_{ij})) &= Pr(\nu_{ij} < A_{ij})E(S_{ij}^* | \nu_{ij} < A_{ij}) \\ &= \frac{1}{\beta(1+\theta_i^c) + \delta(k_i)} Pr(\nu_{ij} < A_{ij})(A_{ij} - E(\nu_{ij} | \nu_{ij} < A_{ij})) \end{aligned} \quad (19)$$

Equation (19) shows how market power affects seller behavior. When the conduct parameter is 0, the slope of the firm's sales with respect to A_{ij} is large, meaning that firms are very responsive to price differentials across regions. By contrast when the conduct parameter is 1, the slope is smaller, meaning that firm sales do not adjust perfectly to pricing differentials, and firms price discriminate in equilibrium by charging higher prices to buyers to whom they have higher spot sales. Thus the elasticity of spot sales with respect to the net price (price minus shipping cost) is what identifies the Cournot model separately from the competitive model. I have used the model to generate many simulated datasets of the same sample size as the dataset I use for estimation, and found that in general I can estimate the conduct parameter quite precisely.

Estimation routine

The parameter vector I have to estimate is $\theta = (\sigma, \phi(.), c_i, \delta(K), \theta_i^c)$. I estimate the model via a maximum likelihood procedure. For each guess of θ, I solve for the policy thresholds A_{ij}, and then use equations (17) and (18) to form the likelihood. I search for the θ that leads to the highest sample likelihood function. Notice that this is essentially a Tobit model where the censoring of spot quantities at zero is explicitly modeled.

The key challenge when estimating the model via maximum likelihood is computing the policy thresholds A_{ij}. For any guess of $\theta = (\sigma, \phi(.), c_i, \delta(K), \theta_i^c)$, I solve the policy functions by the following fixed point algorithm:

- Begin with an initial guess of A_{ij} $\forall i, j$.
- At step $(k+1)$, we know the policy threshold A_{ij}^k from the previous step k. Then for each i, j:
  - Compute $\sum_{i' \neq i} E(S_{i'j}(\nu_{i'j}))$ and $\sum_{k \neq j} E(S_{ik}(\nu_{ik}))$ using A_{ij}^k and equation (19).
  - Update the policy threshold to A_{ij}^{k+1} using equation (15).
- Continue iterating till $||A_{ij}^{k+1} - A_{ij}^k|| < \varepsilon$, where ε is a sufficiently small threshold.

Results

In this section I describe estimates of a relatively parsimonious version of this model. I do not include any covariates in x_{ij} so that the trade costs between i and j are purely given by the (observed) shipping costs plus the random match cost ν_{ij}. Moreover I assume that $\delta(k_i) = \frac{\delta}{k_i}$, meaning that marginal costs are increasing linearly in capacity utilization (q_i/k_i). The parsimonious model leads to a reasonably good fit of prices and exports, and adding more covariates does not measurably change the model fit.

In the main specification (column (2) of Table 4), each buyer is a separate importing country. In an alternative specification, I aggregate buyers into nine importing regions: North East Asia, Southeast Asia, South Asia, Middle East, Southern Europe, Northern Europe, Central America, North America and South America.

Table 4 shows that the estimated conduct parameter is estimated to be 0.72 (when we treat each importing country as a separate market) and 0.90 (when we treat each importing region as a separate market). A conduct parameter of 1 would indicate that firms are behaving exactly as the Cournot model would predict. The conduct parameter estimate of 0.72-0.90 suggests that the Cournot model may be a reasonable description of how sellers choose spot quantities in practice, though it also suggests that seller behavior may be more competitive than Cournot behavior. It appears unlikely that sellers engage in competitive behavior, since the estimated conduct parameter is very far from 0.

The estimated δ parameter suggests that increasing capacity utilization by 10% increases marginal costs by \$1.5/MMBtu, compared to an average price of \$8.6/MMBtu. This indicates that capacity constraints are an important component of marginal costs in this industry.

Table 4: Conduct and Cost Parameter Estimates

	(1)	(2)
SD of match shocks, σ	3.08	3.72
Conduct parameter, θ^c	0.90	0.72
Convexity, δ	7.47	8.59
N	1342	3353
Regions	9	40
Fit (R^2)		
Prices p_j	0.66	0.70
Total Exports by Exporter q_i	0.98	0.98
Total Spot Exports by Exporter $\sum_j S_{ij}$	0.73	0.74
Spot Trade Flows S_{ij}	0.55	0.46

Note: Each observation is an exporter-importer-year pair. Total LNG spot exports from country i to country j in year t are measured in million tonnes. The spot price and shipping costs are measured in \$/MMBtu. In column (1), there are 9 importing markets: North East Asia, Southeast Asia, South Asia, Middle East, Southern Europe, Northern Europe, Central America, North America and South America. In column (2), each importing country is a separate importing market, and there are 40 markets in total.

Figure 9: Actual vs. Predicted Spot Prices

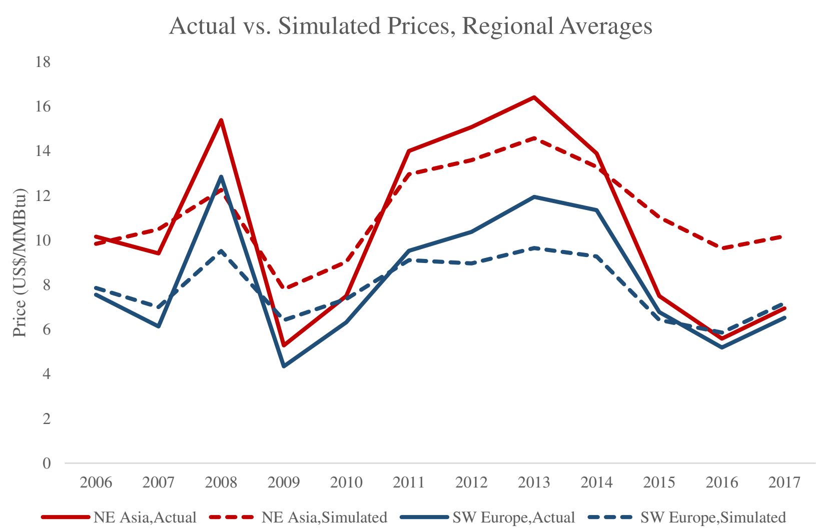

Note: The figure plots actual and predicted annual prices in Northeast Asia (comprising Japan, China, Korea and Taiwan) and South-western Europe (comprising Spain, France, and Italy). The predicted prices are calculated by simulating the estimated model (specification (2) in Table 4) for 1000 different match shocks, and averaging across each of these simulations.

Model Fit

The model fits many of the key data patterns reasonably well, as shown in the last panel of Table 4. The R^2 for prices is 0.66-0.70 (depending on the exact specification), while the R^2 for spot trade flows is 0.46-0.55. The R^2 for total exports is close to 1, but this is because total exports are largely pinned down by contracts. Even so, when we look at total spot exports by seller, the R^2 is reasonably high at around 0.73-0.74.

Figure 9 shows actual versus predicted prices in Northeast Asia and South-western Europe, using specification (2) in Table 4. Though the model does not predict prices precisely, it is able to predict both the small price differentials prior to 2011, and the fact that prices in Asia diverged from prices in Europe starting from 2011 and lasting till 2014. One feature of the data that the model does less well at predicting is the extent of the price convergence from 2015 onward: though the model-predicted price differentials do shrink from 2015 onward, they still remain substantially higher compared to the true price differentials. This period coincided with the building of new capacity in Australia and the US as well as a drop in shipping costs. There are two possible reasons why the model does less well in the last three years of the sample. One is that it treats exporting countries as decision-making entities. This might be a reasonable assumption for countries such as Qatar, Indonesia and Malaysia (where LNG exporting is under the control of a national oil company), but likely to be less tenable for Australia and US (where each terminal operates as a separate decision-making entity), and these are precisely the two countries that expand capacity in the latter part of the sample. A second possibility is that perhaps the conduct parameter is not fixed over time, and the market became more competitive after the entry of new suppliers, so that suppliers acted more competitively than the Cournot model would predict after 2014.

Consequences of market power

We now explore the effects of market power using the estimated structural model. We do so by carrying out a simple counter-factual experiment where we set the conduct parameter to 0. This means that sellers now behave as price-takers, and take expected prices as given when choosing their spot flows to different destinations.

When we shut down sellers' ability to exercise market power, we estimate there is a welfare gain of 12 bn USD (over 12 years). The estimated welfare gain is sizable in relative terms, as it equals 4.5% of total revenue to exporters from the spot market.

Figure 10 shows price differentials across regions in the baseline simulation with market power and the counter-factual simulation without market. We find that price differentials are much smaller when we remove seller ability to price discriminate. This effect is especially pronounced in the period between 2011 and 2013, following the Fukushima disaster in Japan. We find that the price differential between Asia and Europe would average around \$2.5 between 2011 and 2013 if sellers behaved competitively, versus a price differential of around \$4.5 when sellers price discriminate. Thus 45% of the post-Fukushima price differential between Asia and Europe is explained by seller market power in the LNG spot market.

Figure 10: Price Differential between Northeast Asia and Southwestern Europe: Baseline vs. Competitive

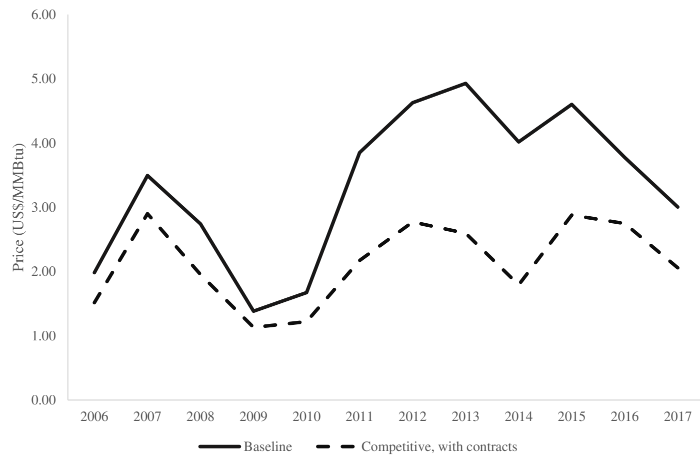

Note: The figure plots the difference between prices in Northeast Asia (comprising Japan, China, Korea and Taiwan) and South-western Europe (comprising Spain, France, and Italy), for both the baseline scenario with Cournot competition and the counter-factual scenario with no market power. The predicted prices are calculated by simulating the estimated model for 1000 different match shocks, and averaging across each of these simulations.

6 Conclusion

I study whether sellers exercise market power in the global liquefied natural gas (LNG) industry, and examine the consequences of market power for pricing, trade and allocative efficiency. Using data from 2006 to 2017 on spot market trade flows, spot prices, shipping costs and seller capacities, I develop and estimate a structural model of LNG trade and pricing that incorporates spatial differentiation, capacity constraints, and trade frictions and flexibly nests different models of seller market power. My method for inferring market conduct builds on the observation that sellers exercising market power engage in third-degree price discrimination and the fact that the Cournot and competitive models have qualitatively different predictions for how capacity-constrained sellers allocate sales across destinations with varying demand shocks. I find that seller decisions are consistent with a Cournot model and unlikely to be generated by a competitive model.

The analysis in this paper has a number of implications. It suggests that buyer concerns about an “Asian price premium" may not be misplaced: without market power, the average price difference between Asia and Europe paid between 2011 and 2013 (the period with the highest price differentials) would have fallen from \$4.5/MMBtu to \$2.5/MMBtu. Thus, market power exacerbates price gaps created by capacity constraints and high shipping costs. The total deadweight loss from market power is also substantial: it is estimated to be USD 12 billion, or about 4.5% of total spot market revenue. More broadly, the method for inferring market power developed in this paper can be applied to other settings where sellers are capacity constrained and their main margin for exercising market power is not by reducing total production but instead by re-allocating their sales across different spatially segmented markets.

References

Asker, John, Allan Collard-Wexler, and Jan De Locker, "(Mis)Allocation, Market Power, and Global Oil Extraction," *American Economic Review*, 2019, *109* (4), 1568–1615.

Boik, Andre, "The empirical effects of competition on third-degree price discrimination in the presence of arbitrage," *Canadian Journal of Economics*, 2017, *50* (4), 1023–1036.

Borenstein, Severin, James B. Bushnell, and Frank A. Wolak, "Measuring Market Inefficiencies in California's Restructured Wholesale Electricity Market," *American Economic Review*, 2002, *92* (5), 1376–1405.

Bresnahan, Timothy F., "The oligopoly solution concept is identified," *Economics Letters*, 1982, *10* (1), 87–92.

_ , "Competition and collusion in the American automobile oligopoly: The 1955 price war," *Journal of Industrial Economics*, 1987, *35* (4), 457–482.

Dafny, Leemore S., "Are health insurance markets competitive?," *American Economic Review*, 2010, *100* (4), 1399–1431.

Fabra, Natalia and Juan Toro, "Price Wars and Collusion in the Spanish Electricity Market," *International Journal of Industrial Organization*, 2005, *23* (3-4), 155–181.

Genesove, David and Wallace P. Mullin, "Testing Static Oligopoly Models: Conduct and Cost in the Sugar Industry, 1890-1914," *RAND Journal of Economics*, 1998, *29* (2), 355–377.

GIIGNL, "The LNG Industry in 2012," 2012. GIIGNL Annual Report.

Grennan, Matthew, "Price discrimination and bargaining: Empirical evidence from medical devices," *American Economic Review*, 2013, *103* (1), 145–177.

Hortacsu, Ali and Steven L. Puller, "Understanding strategic bidding in multi-unit auctions: a case study of the Texas electricity spot market," *RAND Journal of Economics*, 2008, *39* (1), 86–114.

Koyama, Ken, "An Analysis on Asian Premium for LNG Price," 2012. Special Bulletin of The Institute of Energy Economics, Japan.

Li, Hui, "Intertemporal Price Discrimination with Complementary Products: E-Books and E-Readers," 2015. SSRN Working Paper.

Miller, Nathan and Matthew C. Weinberg, "Understanding The Price Effects Of The Miller-Coors Joint Venture," *Econometrica*, 2017, *85* (6), 1763–1791.

_ **and Matthew Osborne**, "Spatial differentiation and price discrimination in the cement industry:

evidence from a structural model," *RAND Journal of Economics*, 2014, *45* (2), 221–247.

Nevo, Aviv, "Measuring Market Power in the Ready-to-Eat Cereals Industry," *Econometrica*, 2001, *69* (2), 307–342.

Puller, Steven L., "Pricing and Firm Conduct in California's Electricity Market," *The Review of Economics and Statistics*, 2007, *89* (1), 75–87.

Ritz, Robert, "Price discrimination and limits to arbitrage: An analysis of global LNG markets," *Energy Economics*, 2014, *45* (1), 324–332.

Salvo, Alberto, "Inferring market power under the threat of entry: the case of the Brazilian cement industry," *RAND Journal of Economics*, 2010, *41* (2), 326–350.

Sweeting, Andrew, "Market Power In The England And Wales Wholesale Electricity Market 1995-2000," *Economic Journal*, 2007, *117* (1), 654–685.

Villas-Boas, Sofia B., "An Empirical Investigation of the Welfare Effects of Banning Wholesale Price Discrimination," *RAND Journal of Economics*, 2009, *40* (1), 20–46.

Williamson, Oliver E., "Transaction-Cost Economics: The Governance of Contractual Relations," *Journal of Law and Economics*, 1979, *22* (2), 233–61.

Wolak, Frank A., "An Empirical Analysis of the Impact of Hedge Contracts on Bidding Behavior in a Competitive Electricity Market," *International Economic Journal*, 2000, *14* (2), 1–40.

Appendix

A Monte Carlo Simulations

I describe Monte Carlo simulations of a simplified version of the model developed in Section 4. In the simplified version of the model, marginal costs are linear until production hits a binding capacity constraint. Trade costs are deterministic and only a function of the distance between countries. The purpose of the simulations is to illustrate the kinds of trade flows and prices generated by Cournot competition and contrast these with the trade flows and prices generated by competitive behavior.

A.1 Assumptions

Throughout the simulations I assume there are 2 sellers and 2 buyers. Seller 1 is located closer to Buyer 1 than Seller 2, and Seller 2 is located closer to Buyer 2 than Seller 1. I assume that firms do not sign any contracts (so that all trade happens using the spot market). I also set re-sale costs very high so that buyers are not able to engage in re-exports in any equilibria; this is done for simplicity so as to focus on seller decisions.

I assume that firms cannot produce beyond their capacity k_i and have constant marginal costs for any output below k_i. I set the marginal cost of production $c_1, c_2 = 4$. I assume that demand shocks in the two buying countries are uniformly and independently distributed: $\varepsilon_j \sim U[d_l, d_h]$, $j = 1, 2$.

Table 5: Monte Carlo Assumptions

Demand Slope	$\beta = 1$	
Demand Shock	$\varepsilon_j \sim U[d_l, d_h]$	
Demand Parameters	$d_l = 40$	$d_h = 80$
Marginal cost of production	$c_1 = 4$	$c_2 = 4$
Shipping costs from seller 1	$d_{11} = 1$	$d_{12} = 5$
Shipping costs from seller 2	$d_{21} = 5$	$d_{22} = 1$
Re-sale costs	$\kappa = \infty$	
	Scenarios	
	S1	S2
Description:	High capacity	Limited capacity
K_1	100	15
K_2	100	15

I consider the following scenarios:

1. S1, where firms have lots of capacity relative to demand. The numbers are chosen so that for every possible pair of demand shocks, firms will produce below capacity.

2. S2, where firms have very limited capacity. The numbers are chosen so that for every possible pair of demand shocks, firms will produce at full capacity even when they have no contracts.

I present some comparisons of the Cournot model and the perfectly competitive model. I look at both scenario S1 (where firms have lots of capacity) and scenario S2 (where firms have very limited capacity). For each scenario, I draw 1000 random demand shocks ε_1 and ε_2 from the uniform distribution and then solve for the perfectly competitive outcome and the Cournot outcome. Table 6 shows average prices, trade flows, production and welfare across all 1000 demand shock realizations.

When firms have lots of capacity (scenario S1), firms withhold production under Cournot competitive, relative to the competitive allocation. This leads to higher prices and lower welfare under Cournot. Another feature of Cournot competition is that sellers are more likely to sell to faraway buyers. Under perfect competition, seller 1 only sells to buyer 1 and seller 2 only to buyer 2. Under Cournot, seller 1 often sells to buyer 2 and seller 2 to buyer 1. This happens because the nearby seller withholds quantity from the nearest buyer, which sometimes makes it profitable for the faraway seller to enter the market.

When firms have very limited capacity (scenario S2), total production is pinned down by capacity and does not depend on the type of competition. However welfare is still lower under Cournot competition. There are two sources of this inefficiency. One is that transportation costs are not minimized under Cournot since buyers are more likely to buy from far away. For example seller 1 sells 83% of its output to the nearest buyer under perfect competition but only 62% under Cournot. This happens even in scenarios with similar demand in the two markets. An example is given by Table 7 which shows that even when demand is very similar in the two markets, the Cournot outcome is worse from a welfare perspective because there is too much trade happening between faraway seller-buyer pairs.

A second source of inefficiency from Cournot is that when demand is higher in one market than other, oligopolists withhold quantity from the high-demand market and over-supply the low-demand market. This is shown in Table 8, which looks at what happens when demand is very high in market 1 but very low in market 2. The competitive allocation involves selling all the gas to market 1; under Cournot, however, seller 2 withholds some of its output from market 1 and sells to market 2 instead, raising prices in market 1 and reducing them in market 2. As such, price differentials across regions are larger under Cournot than in perfect competition.

The Cournot model generates the following stylized facts. First, there is imperfect sorting of trade flows by distance: sellers often sell LNG to buyers that are very far away, despite the presence of nearby buyers. Second, sellers engage in price discrimination: they earn higher markups from some buyers than other buyers. In the competitive model there is no price discrimination, and less likelihood that sellers will sell to faraway buyers. These are all consistent with LNG trade flows and

Table 6: Comparison of competitive and Cournot outcome,

		Prices		Trade flows					
				Seller 1		Seller 2			
	Welfare	p_1	p_2	q_{11}	q_{12}	q_{21}	q_{22}	Total prod.	Total capacity
S1: High capacity									
Competitive	3152	5	5	55	0	0	55	110	200
Cournot	2626	25	25	20	16	16	20	71	200
S2: Limited capacity									
Competitive	1466	44.8	45.1	12.5	2.5	2.2	12.8	30	30
Cournot	1438	44.8	45.1	9.4	5.6	5.4	9.6	30	30

Note: For each scenario, I draw 1000 different pairs of demand shocks ($\varepsilon_1, \varepsilon_2$). For each realization of the demand shock, I solve for the competitive spot prices and allocation as well as the Cournot prices and allocation. The table presents averages across these 1000 realizations. For example, p_1 refers to the average price paid by buyer 1 across all realizations of demand shocks.

Table 7: Scenario with moderate demand in both markets

ε_1	ε_2		Welfare	Price in market 1	Price in market 2	q_{11}	q_{12}	q_{21}	q_{22}
55.9	56.7	Competitive	1314	40.9	41.7	15	0	0	15
		Cournot	1270	41.2	41.5	9.4	5.6	5.4	9.6

Table 8: Scenario with high demand in market 1, low demand in market 2

ε_1	ε_2		Welfare	Price in market 1	Price in market 2	q_{11}	q_{12}	q_{21}	q_{22}
78.3	41.0	Competitive	1689	48.3	41.0	15	0	15	0
		Cournot	1671	51.2	38.1	15	0	12.1	2.9

prices that are observed in practice. It is also instructive to note the predictions which are similar for the Cournot and competitive models. Both models produce price dispersion and convergence: prices are different across regions during periods when demand is different, but prices are similar when demand shocks are similar (though the Cournot model predicts higher price differentials during periods with asymmetric demand). The Cournot model with capacity constraints also does not predict *average* prices that are higher than perfect competition; thus the spatial nature of trade flows is essential in order to tell apart Cournot competition from perfect competition.

B Derivation of first-order condition in the case of linear demand and quadratic costs

In the case of linear demand and quadratic costs, the first-order condition (10) becomes linear in the match cost ν_{ij}, as we show below.

Because of linear demand, expected prices and the derivative of expected prices with respect to S_{ij} are given by the following equations:

$$\begin{aligned} E(p_j^*|S_{ij}(\nu_{ij})) &= \beta(\varepsilon_j - Q_j^c - S_{ij} - \sum_{i' \neq i} E(S_{i'j}(\nu_{i'j}))) \\ \frac{\partial E(p_j^*|S_{ij}(\nu_{ij}))}{\partial S_{ij}} &= -\beta \end{aligned}$$

Because of the assumption of quadratic costs of production (see equation (13)), we can write expected costs and the derivative of expected costs with respect to S_{ij} as:

$$\begin{aligned} C(q_i, k_i) &= c_i q_i + \frac{\delta(k_i)}{2} q_i^2 \\ E(C(q_i, K_i)|S_{ij}(\nu_{ij})) &= c_i\big(q_i^c + S_{ij}(\nu_{ij}) + \sum_{k \neq j} E(S_{ik}(\nu_{ik}))\big) + \frac{\delta(k_i)}{2}[q_i^c + S_{ij}(\nu_{ij})]^2 \\ &\quad + 2\frac{\delta(k_i)}{2}[q_i^c + S_{ij}(\nu_{ij})] \sum_{k \neq j} E(S_{ik}(\nu_{ik})) + \frac{\delta(k_i)}{2} E\big(\sum_{k \neq j} S_{ik}(\nu_{ik})\big)^2 \\ \frac{\partial E(C(q_i, k_i)|S_{ij}(\nu_{ij}))}{\partial S_{ij}} &= c_i + \delta(k_i)\big(q_i^c + S_{ij}(\nu_{ij}) + \sum_{k \neq j} E(S_{ik}(\nu_{ik}))\big) \end{aligned}$$

Plugging these expressions into equation (8), we get the following equation:

$$\begin{aligned} &\beta(\varepsilon_j - Q_j^c - S_{ij} - \sum_{i' \neq i} E(S_{i'j}(\nu_{i'j}))) - \theta_i^c \beta S_{ij}(\nu_{ij}) - \\ &\big(c_i + \delta(k_i)\big(q_i^c + S_{ij}(\nu_{ij}) + \sum_{k \neq j} E(S_{ik}(\nu_{ik}))\big)\big) - \big(d_{ij} + \phi(x_{ij}) + \nu_{ij}\big) \leq 0 \end{aligned}$$

Collecting terms:

$$\begin{aligned} &\beta\big(\varepsilon_j - Q_j^c - \sum_{i' \neq i} E(S_{i'j}(\nu_{i'j}))\big) - \big(c_i + \delta(k_i)\big(q_i^c + \sum_{k \neq j} E(S_{ik}(\nu_{ik}))\big) + d_{ij} + \phi(x_{ij})\big) - \\ &\qquad \big(\beta(1 + \theta_i^c) + \delta(k_i)\big) S_{ij}(\nu_{ij}) - \nu_{ij} \leq 0 \qquad (20) \end{aligned}$$