Queen's Economics Department Working Paper No. 1518

Long-term contracts and efficiency in the liquefied natural gas industry

Nahim Bin Zahur

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

1-2024

Long-term contracts and efficiency in the liquefied natural gas industry

Nahim Bin Zahur *

January 2024

Abstract

In many capital-intensive markets, sellers sign long-term contracts with buyers before committing to sunk cost investments. Ex-ante contracts mitigate the risk of under-investment arising from ex-post bargaining. However, contractual rigidities reduce the ability of firms to respond flexibly to demand shocks. This paper provides an empirical analysis of this trade-off, focusing on the liquefied natural gas (LNG) industry, where long-term contracts account for over 70% of trade. I develop a model of contracting, investment and spot trade that incorporates bargaining frictions and contractual rigidities. I structurally estimate this model using a rich dataset of the LNG industry, employing a novel estimation strategy that utilizes the timing of contracting and investment decisions to infer bargaining power. I find that without long-term contracts, sellers would decrease investment by 27%, but allocative efficiency would significantly improve. Negative contracting externalities lead to inefficient over-use of long-term contracts in equilibrium. Policies aimed at eliminating contractual rigidities reduce investment by 16%, but raise welfare by 9%.

Keywords: Long-term Contracts, Spot Markets, Under-investment, Nash Bargaining, Contracting Externalities, Market Power, Liquefied Natural Gas

*Department of Economics, Queen's University (email: nahim.zahur@queensu.ca). I am especially grateful to my advisors Panle Jia Barwick, Jean-Francois Houde, Shanjun Li, and Giulia Brancaccio for invaluable advice and guidance, and to Robert Clark and Ricard Gil for helpful comments and suggestions. I would like to thank Abhi Ananth, Matt Backus, Meredith Fowlie, Todd Gerarden, Matthew Gentzkow, Myrto Kalouptsidi, Ryan Kellogg, Alex MacKay, Francesca Molinari, Mar Reguant, Ivan Rudik, Jorg Stoye, Lucy Xiaolu Wang, and seminar participants at the Bank of Canada, Dallas Fed, Michigan, Penn State, Queen's, Waterloo, ASSA Meetings, Interactive Online IO Seminar, IIOC, Econometric Society Meetings, EEA-ESEM, IO Canada, SEA, WEAI, NE Workshop on Energy Policy and Environmental Economics & Energy Research Insights for Decision-making Conference 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

In many business-to-business markets, sellers make large sunk cost investments before production begins. They sell to buyers that may have sizeable bargaining power, because of limited availability of alternative buyers, relationship-specific investments, or search frictions. Examples include markets for automobile parts (Klein et al., 1978), coal (Joskow, 1987), electricity (Bushnell et al., 2008; Ryan, 2021), iron ore (Wilson, 2012), and trucking (Hubbard, 2001).

A key question that arises in such markets is whether long-term contracts enhance or hinder economic efficiency, especially when firms also have the option of trading on spot markets. A benefit of long-term contracts is that they facilitate sunk cost investments (Williamson, 1975; Grossman and Hart, 1986; Edlin and Reichelstein, 1996). If agents negotiate the terms of trade after the investment is sunk, sellers may be unable to recoup the full investment cost in ex-post bargaining, leading to under-investment. Negotiating long-term contracts *ex-ante* (prior to investment) mitigates the risk of under-investment. But a drawback of long-term contracts is that they may be inflexible in response to fluctuations in demand and costs, as it is costly to account for all possible contingencies when writing a contract (Masten and Crocker, 1985). By contrast, trading on spot markets allows firms greater flexibility in adjusting to uncertainty, potentially leading to greater allocative efficiency.

The trade-off between under-investment and contract inflexibility is particularly salient in the global liquefied natural gas (LNG) industry, among the fastest-growing energy markets in the world. LNG sellers make large upfront investments in liquefaction terminals that convert natural gas to LNG. Long-term contracts are valuable for a variety of reasons (such as supply assurance), but buyers have considerable bargaining leverage in contract negotiations: there is only a limited pool of buyers available to contract with at any point in time, and switching to a different buyer is costly due to shipping costs. Under-investment is therefore a natural concern, and sellers typically only invest after first signing ex-ante long-term contracts with buyers (on average 20 years in length, and accounting for 60% of their capacity). At the same time, the LNG market is subject to large demand fluctuations, such as the Fukushima nuclear disaster in 2011 (which raised Japan's LNG demand) and the reduction in Russian natural gas exports to Europe in 2022 (which raised demand for LNG among European buyers). Long-term LNG contracts are potentially inflexible in responding to such demand shocks, especially as they often include clauses that explicitly prohibit resales.[1]

In this setting, I answer two research questions. First, what are the costs and benefits of long-term contracting? Second, are there welfare gains from regulating long-term contracting? Within the LNG industry, there has been considerable debate over whether long-term contracts should be regulated. Anti-trust authorities in the European Union and Japan have introduced regulations

[1] These are known as "destination clauses" and prohibit buyers from re-selling LNG outside their home country.

limiting the use of resale restrictions in LNG contracts, in an attempt to address the inflexibility of long-term contracts, but there is a concern that such regulations may reduce sellers' incentives to invest (by limiting their ability to exercise market power). Assessing the welfare effects of these regulations therefore requires quantifying their effects on both allocations and investment.

To answer these questions, this paper estimates a structural model of the LNG industry. I develop a multi-stage game where sellers make investments and sellers and buyers negotiate long-term contracts, with the outside option of trading on the spot market. I propose a novel estimation strategy that leverages the timing of contracting and investment decisions to distinguish between under-investment and other motives for contracting (such as supply security). This approach allows the researcher to learn about bargaining power in the absence of data on contract prices, a common data challenge when studying long-term contracts. Using the estimated model, I study the consequences of long-term LNG contracts for investment and allocative efficiency, and assess the welfare effects of regulating long-term contracts.

The model features spatially differentiated sellers and buyers. Buyers have stochastic demand for LNG which cannot be perfectly predicted in advance. I allow buyers to have a different willingness-to-pay for contracted and spot LNG, since buyers may prefer long-term contracts due to supply assurance motives (Bolton and Whinston, 1993), and reduced transaction costs and search frictions (MacKay, 2022; Tolvanen et al., 2022). Sellers incur fixed costs of investing in liquefaction terminals, as well as variable costs of producing and shipping LNG in every period.

The sellers and buyers play a sequential, multi-stage game. In the first stage, they negotiate *ex-ante* long-term contracts, before the seller has invested. I assume the equilibrium outcome of these negotiations is described by the “Nash-in-Nash" bargaining solution: each seller-buyer pair Nash bargains over the contract quantity and a lump-sum transfer to be paid by the buyer to the seller, taking as given the contracts agreed to by all other buyers and sellers (Horn and Wolinsky, 1988; Chipty and Snyder, 1999). In the second stage, the seller chooses how much to invest and pays the sunk cost of investment. The seller must build enough capacity to satisfy ex-ante contracts signed in the first stage. In the third stage, after the seller has committed to the investment, sellers and buyers negotiate *ex-post* contracts (once again via Nash bargaining). Finally, every year, demand shocks are realized, and the sellers and buyers participate in the global LNG spot market, where sellers sell surplus LNG not already committed under long-term contracts. Sellers engage in Cournot competition on the spot market. In equilibrium this results in spatial price discrimination, with buyers in different regions paying different spot prices.

The model features the key economic mechanisms underlying firms' choices over how to trade. First, sellers under-invest if they were to only rely on ex-post contracts and if buyers have bargaining leverage. Signing ex-ante long-term contracts allows the seller and buyer to mitigate under-investment. The risk of under-investment (and the resulting incentive to sign larger ex-ante contracts) is larger if the seller's bargaining leverage is weaker relative to the buyer, which is a function of the value of their outside options of trading on the spot market. For instance, the risk of under-investment is larger when the seller is located far away from alternative buyers and must incur high shipping costs to trade with them. Conversely, the risk of under-investment is smaller if sellers have more market power on the spot market, since market power strengthens the seller's outside option and weakens the buyer's outside option.

Second, long-term contracts lock in transactions before the parties have full information about demand. This inflexibility is costly as it restricts the ability of capacity-constrained sellers to meet demand shocks by reallocating LNG across buyers. But because sellers exercise market power on the spot market, long-term contracts also have pro-competitive effects (Allaz and Vila, 1993), since they fix in advance the price sellers receive for a portion of their sales, and thus reduce their incentives to withhold production as a way to push up spot prices. The allocative efficiency consequences of using long-term contracts depends on which of these two effects dominates.

Third, a long-term contract negotiated between a seller and a buyer imposes negative externalities on other buyers (Bolton and Whinston, 1993; Segal, 1999). This is because contracts commit part of the seller's output for the exclusive use of one buyer, reducing the quantity of LNG available to other buyers competing for limited LNG supplies. Due to contracting externalities, the equilibrium level of contracting is in general not socially optimal.

The empirical analysis uses a rich dataset on capacity, long-term contracts, trade flows, and spot prices in the LNG industry collected from various industry sources, spanning a period from 2004 to 2017. The dataset includes the universe of LNG investments and long-term contracts operational during this period. Although the negotiated contract price is unobserved, I observe the contract duration, quantity and signature date for each contract, as well as the date of "final investment decision" (i.e., the date when the seller commits to building the terminal), start date, and capacity of each investment. The ability to observe the precise timing of contract signatures and investment decisions is a unique feature of the data and is crucial for the identification of the structural model.

Several descriptive data patterns suggest that concerns about potential under-investment are an important determinant of contracting behavior. First, contracting primarily takes place before investment: on average, a seller signs ex-ante contracts amounting to around 60% of their capacity before making a final investment decision. Second, sellers only have access to a limited number of buyers with whom they can sign long-term contracts at any point in time: in an average year, 12 buyers sign long-term contracts.[2] Third, I construct a measure of the strength of the outside option of each negotiating party, by computing how far away they are from alternative trading partners (relative to their chosen trading partner). I find, consistent with the predictions of the theory,

[2]These contracts also differ considerably in size, so that switching to a different buyer is likely to entail a non-trivial adjustment in the contract quantity (making it costly for the seller to switch to a different buyer).

that firms sign larger ex-ante contracts when sellers have weaker outside options, but sign smaller ex-ante contracts when buyers have weaker outside options.

I estimate the structural model in several steps. I first estimate demand curves for each buyer using demand shifters in other markets (that are excluded from their own demand) as instruments for the spot price. I then estimate seller production costs by exploiting the first-order conditions of the Cournot game and data on spot trade flows, spot prices and shipping costs. These estimates are used to construct expected payoff functions for sellers and buyers, which correspond to their disagreement payoffs (or outside options) during the contract negotiations.

Next, I estimate a set of parameters characterizing the contracting and investment decisions: the investment cost, buyer preferences for contracting, and a bargaining weight parameter that governs the distribution of surplus between seller and buyer. The estimation utilizes the equilibrium conditions of the bargaining and investment game, solved via backward induction. A key challenge in identifying the bargaining weight is that long-term contract prices (which would be the natural source of information on how sellers and buyers split the surplus from trade) are unobserved. To overcome this challenge, I leverage variation in the outside options of sellers and buyers across different negotiations and over time. Intuitively, the extent to which firms adjust the size of ex-ante and ex-post contracts and the size of investments, in response to variation in the outside options of sellers and buyers, identifies the bargaining weight. For instance, if the seller's bargaining weight is high (meaning that sellers capture most of the surplus from bargaining), changes in the seller's outside option will have little effect on the negotiated price, and therefore have little effect on the equilibrium sizes of contracts and investments.

I find that sellers are highly capacity-constrained in the short-run and face large sunk costs of investment in the long-run, with an average export terminal estimated to cost $22 bn to build. The seller bargaining weight is estimated to be 0.64, and the hypothesis that sellers make take-it-or-leave-it offers is rejected by the data. Sellers' incentives to invest are thus dampened by buyer bargaining power: the marginal benefit from investing would be 22% higher if sellers were able to fully enjoy the surplus from investing. The potential for under-investment creates incentives to sign large *ex-ante* long-term contracts, despite the fact that (investment effects aside) buyers ideally would prefer to trade via *ex-post* contracts: they are willing to pay a 15% premium to trade using ex-post contracts rather than spot, but only a 1% premium to trade using ex-ante contracts. Finally, negative contracting externalities are significant: a long-term contract imposes a marginal external cost on other agents equal to $0.9/MMBtu on average (10% of the average spot price).[3]

Using the estimated model, I carry out various counter-factual exercises. First, I evaluate the trade-off between under-investment and contract inflexibility. I begin by quantifying the role of

[3]As a validation of the methodology of estimating bargaining power without observing negotiated prices, I compare the contract prices predicted by the model with contract prices that can be inferred from customs data for a subset of the contracts. The model-predicted contract prices match up well with these external measures of contract prices.

long-term contracts in mitigating under-investment. I find that if sellers are not able to sign ex-ante or ex-post long-term contracts with buyers, they would lower investment by 27%. The reduction in investment is primarily due to the inability to sign *ex-ante* contracts: if firms can sign ex-post but not ex-ante contracts, investment still decreases by 24%. Sellers who are geographically more isolated (i.e., located farther away from their nearest buyers than the median seller) reduce investment by 30% when they cannot sign ex-ante contracts, compared to 11% for sellers located closer to buyers.

Next, I quantify the allocative efficiency consequences of using long-term contracts, holding capacity fixed. I find that switching from long-term contracts to spot trade would result in allocative efficiency gains of $38 bn. While removing long-term contracts worsens the deadweight loss from market power on the spot market (à la Allaz and Vila, 1993), this is outweighed by the flexibility gains from greater use of the spot market. These flexibility gains arise because firms are more efficient at responding to demand shocks when their capacity is not tied up under long-term contracts. To further investigate this, I consider the effect of a shutdown of Russian natural gas exports to Europe, leading to a large increase in European demand for LNG. I simulate the industry response to this demand shock both with and without long-term contracts. I find long-term contracts result in a more muted response to the demand shock than is efficient, since some sellers bound by long-term contracts do not re-allocate LNG to Europe.[4]

Second, I evaluate the efficiency of long-term contracting, by simulating the industry equilibrium with and without long-term contracts. Eliminating long-term contracting results in welfare gains equal to $23 bn, or about 0.4% of total welfare: the reduced investment by sellers (when they cannot contract) is more than compensated for by gains in allocative efficiency. Contracting externalities lead to an inefficiently high degree of long-term contracting in equilibrium, explaining why switching from long-term contracts to spot can improve efficiency (despite the value of contracts in curbing under-investment).

Third, I assess the welfare effects of a policy banning the use of resale restrictions in long-term contracts. The policy reduces the ability of sellers to engage in spatial price discrimination (since they face a stronger threat of arbitrage), and thus weakens their outside option from trading on the spot market. As such, sellers reduce investment by 16%. Despite the reduction in investment, the removal of resale restrictions leads to a substantially more efficient allocation of LNG, leading to sizeable welfare gains of $513 bn (or 9.4%).

Contributions to the Literature: This paper contributes to three main strands of literature. First, it builds on an extensive theoretical (Williamson, 1975; Klein et al., 1978; Grossman and Hart, 1986; Edlin and Reichelstein, 1996) and empirical (Joskow, 1987; Crocker and Masten, 1988;

[4]Consistent with this, US exporters in July 2022 cited existing contractual commitments as a constraint on their ability to meet increased European demand for LNG: see `https://www.wsj.com/articles/worlds-growing-thirst-for-american-gas-tests-u-s-ability-to-meet-demand-11658494858`.

Hubbard, 2001) literature on long-term contracting.[5] The paper is most closely related to a recent literature that uses structural models to study the costs and benefits of long-term contracts (Tolvanen et al., 2022; MacKay, 2022) and the efficiency consequences of weak contract enforcement (Ryan, 2020) and hold-up risk (Bhattacharya, 2021; Ryan, 2021). The main contribution of this paper is to provide a novel empirical framework for quantifying the trade-off between under-investment and inflexibility that firms face when using long-term contracts. In contrast to a recent literature that has emphasized inefficiencies from weak contract enforcement (for example, Nunn, 2007, Blouin and Macchiavello, 2019 and Ryan, 2021), this paper highlights how long-term contracts can result in inefficiencies even when they are enforceable.

The paper contributes to a growing literature in industrial organization that uses bargaining models to study negotiations between firms, based on the Horn and Wolinsky (1988) framework. Most papers in this literature infer bargaining power from negotiated prices (e.g., Crawford and Yurukoglu, 2012; Grennan, 2013; Gowrisankaran et al., 2015; Ho and Lee, 2017). My contribution is to provide a new strategy for inferring bargaining power from the timing and size distribution of contracts and investment, using the insight that firms with lower bargaining leverage are more likely to under-invest and therefore have a stronger incentive to sign large ex-ante contracts.[6] This strategy can be useful in other settings where the researcher lacks data on negotiated prices, but observes contracting and investment decisions that are functions of these negotiated prices.

Finally, the paper contributes to the literature measuring the effects of contracting on short-run allocative efficiency in energy markets. Building on the theoretical insights from Allaz and Vila (1993), a series of papers have empirically analyzed electricity markets and have found that forward contracts significantly mitigate the deadweight loss from seller market power (Bushnell et al., 2008; Ito and Reguant, 2016). In contrast to this literature, I find that allocative efficiency can decrease from the use of long-term contracts, despite the pro-competitive effect of contracts (which also exists in my setting), due to their inflexibility in responding to demand fluctuations. In addition, the paper is also related to a literature studying long-term contracting in the LNG industry (Ruester, 2009; Hartley, 2015; Agerton, 2017). My paper builds on this literature by estimating a new structural model of the LNG industry that can be used to empirically quantify the welfare effects of using long-term contracts.

The rest of the paper is divided as follows. Section 2 discusses key institutional features of the LNG industry and describes the dataset. Section 3 discusses descriptive evidence. Section 4 presents the model. Section 5 describes estimation. Section 6 presents results from counter-factual simulations. Section 7 concludes.

[5] A related literature studies the role of informal long-term relationships and relational contracts: see, for example, Macchiavello and Morjaria (2015); Macchiavello and Miquel-Florensa (2017) and Harris and Nguyen (2022).

[6] The closest antecedent to this strategy that I am aware of was developed by Bhattacharya (2021), who exploits information on ex-ante investments to identify a bargaining parameter.

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 and importing country need to invest in specialized infrastructure. In the exporting country, a liquefaction terminal converts natural gas into LNG, and loads it onto a tanker. The tanker transports LNG to a regasification terminal at the importing country, where the LNG is unloaded, converted back into gaseous form and used for power generation and for heating.

Many countries rely on 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. Japan is the largest importer of LNG, followed by China, Korea, India, Taiwan and Spain. Major exporters of LNG are Australia, Qatar, USA, Malaysia and Indonesia.

The LNG industry has grown rapidly in recent years. Between 2004 and 2017, the number of LNG importing countries increased from 14 to 40, while the number of exporting countries increased from 12 to 19. The total volume of LNG trade more than doubled between 2004 and 2017. By 2021, the value of global annual trade in LNG exceeded US$150 billion.[7]

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. 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 used to determine the price under which trade takes place. The contract price is usually indexed to the price of some benchmark (e.g., the oil price).

In addition to these basic features, a contract may specify a “take-or-pay" quantity (no higher than the average annual quantity): this is a minimum quantity of LNG the buyer commits to paying for every year, whether or not they actually take delivery. This ensures that the seller’s minimum revenue in any given year is the price multiplied with the take-or-pay quantity. A contract may include a destination clause, which prohibit buyers from reselling the product outside a pre-defined market (typically the buyer’s home country) (IEA, 2013). 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 (known in industry parlance as a “diversion").

Both destination clauses, and to a lesser extent diversion clauses, have been controversial in the LNG industry. Destination clauses directly prevent buyers from engaging in arbitrage; as such, they have been challenged by anti-trust authorities in LNG importing countries, though with mixed success. The European Commission ruled destination clauses anti-competitive as far back as 2003

[7]See https://www.canadianenergycentre.ca/wp-content/uploads/2022/06/CEC-Research-Brief-22-V5-June-21-2022.pdf.

(IEA, 2013). Since then, the Commission has successfully negotiated their exclusion from LNG contracts signed with Nigeria and Algeria (Talus, 2011), though some LNG contracts signed by European buyers may still have destination clauses.[8] In 2017, the Japanese Fair Trade Commission (FTC) prohibited destination clauses in the majority of new LNG contracts (Harding and Sheppard, 2017). Diversion clauses, while not directly prohibited by either the European Commission or Japan's FTC, have also come under anti-trust scrutiny as they can reduce the incentives of buyers to engage in arbitrage, since the profits from arbitrage now have to be shared with the seller.[9]

Figure 1 shows the distribution of contract duration for long-term contracts (defined as contracts that are longer than four 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 to trade LNG, however. Parties can also trade LNG using short-term contracts, or on the spot market.[10] Figure 2 shows that the share of spot and short-term trade has been rising over time, from 12% in 2004 to 27% 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, though this accounts for a very small proportion of overall trade.[11]

Appendix Figure A1 shows the evolution of total liquefaction capacity, regasification capacity and LNG trade over time. The binding constraint on the volume of trade tends to be the available liquefaction capacity, and export capacity utilization is high, ranging from between 80 to 90%. By contrast, there is generally a lot of excess regasification capacity. These patterns reflect the fact that liquefaction projects are typically significantly more costly than regasification projects: the capital cost of liquefaction projects is on average equal to $1.7 billion per mtpa (million tonnes per annum) of capacity, compared to $250 million per mtpa for regasification projects (OIES, 2017).

Data

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

Data on LNG contracts: Data on individual LNG contracts were originally collected from Bloomberg. I combine these data with annual industry reports provided by the GIIGNL (The International Group of Liquefied Natural Gas Importers), as well as a dataset of long-term natural gas contracts published by Neumann et al. (2015). In addition, for the vast majority of contracts, I have collected press releases, company reports and newspaper articles announcing the signing of

[8]A 2017 report commissioned by the Commission mentioned that destination clauses were "in the process of being removed" from European contracts (European Commission, 2017).

[9]Talus (2011) distinguishes between two kinds of diversion clauses: "profit-splitting" clauses that split the eventual profits from resale between the buyer and seller, and "price-splitting" clauses that split the resale price. He argues that "price-splitting" clauses often entirely remove the buyer's incentive to resell.

[10]A typical spot transaction involves delivery of a single LNG cargo from the seller's terminal to the buyer's terminal.

[11]The share of re-exports increased from 0.15% in 2008 to a peak of 2.69% in 2014, but decreased to 1% by 2017.

Figure 1: Histogram of long-term contract durations

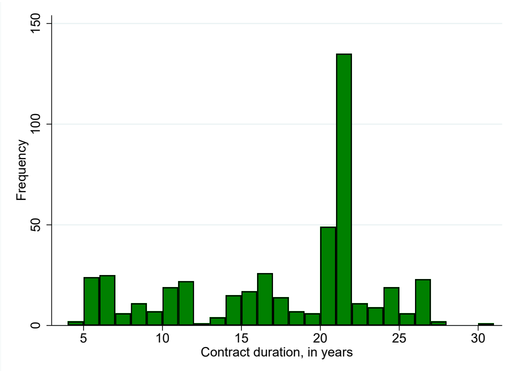

Figure 2: Share of short-term contracts and spot transactions in LNG trade

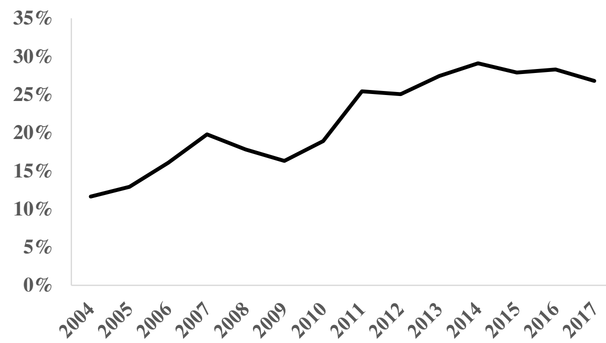

Note: Long-term contracts are contracts exceeding four years in duration (Figure 1). Figure 2 plots trade carried out using short-term contracts (no longer than four years) or on the spot, as a share of total LNG trade. Source: GIIGNL.

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 original dataset. The eventual dataset consists of every long-term LNG contract that was signed in this industry from 2004 to 2017, as well as any long-term contracts signed prior to 2004 that were still active in 2004.

For each contract, I observe the contract quantity, the year when the contract was signed, and the contract start and end year. I observe the identity of the buyer and seller and if the contract is an extension of an earlier, expiring contract. There are a number of contract details that I do not observe. The most important of these is the pricing formula, which is typically confidential and known only to the parties that are signatory to the contract. I also do not observe the "take-or-pay" share of the contract, and whether the contract includes a destination or diversion clause.

Data on liquefaction and regasification capacity: I obtain data on investment and capacity, for both the liquefaction and regasification terminals, from the annual reports of the GIIGNL. The dataset includes the start-up year, nameplate capacity, the ownership structure and the operator for every terminal operational from 2004 to 2017. I complement this with information on the year when a Final Investment Decision (FID) is made on an export terminal, which I hand-collected for every project from official press releases and news articles.[12]

Data on LNG trade flows, spot prices and shipping costs: I utilize two datasets recording LNG trade flows. The first dataset (from Bloomberg) records quarterly LNG trade flows for each country pair from 2003Q4 to 2018Q2. I utilize this dataset for demand estimation. Second, I collect data on LNG trade flows from the annual reports of the GIIGNL. This dataset reports yearly

[12]The FID is the decision by all project partners to finally commit to the project. Construction of an LNG export terminal only begins once a FID has been taken by the investors.

LNG trade flows for each country pair from 2004 to 2017, broken down into trade flows that take place under long-term contracts, short-term contracts and spot trade, and re-exports.[13] Since this dataset distinguishes between flows that take place under long-term contracts and short-term/spot contracts, I use this dataset when estimating the spot trade model. Finally, I obtain data on weekly LNG spot prices and shipping costs between February 2006 and August 2018 from several sources; Appendix A provides more details.

Appendix Table A1 contains summary statistics on key variables used in the analysis. As Panel C shows, the average long-term contract is 17 years in duration, and is signed 3.6 years before the start date of deliveries. Export projects are generally very large in size (Panel D), with the typical investment equal to 6.94 mtpa (for context, the export capacity of an average LNG exporting country is 14 mtpa). Time-to-build is substantial: on average 4.3 years pass between the date of the FID and the time when the export project begins operating.

Table 1: Contract Types

	Contracts with fixed origin and destination		Flexible contracts
	Ex-ante	Ex-post	
Obs.	124	246	94
Annual contract quantity (mtpa)	1.58	1.11	1.34
Duration (years)	20.67	15.74	16.61
Total contract quantity (mt)	32.59	17.25	22.90
Time from signature to start (years)	5.27	2.84	3.44

[1.] The contract quantity, duration and time from signature to start are sample averages for each type of contract.

The majority of long-term contracts specify a fixed export terminal (origin) and a fixed destination, but more recently there has been increasing use of flexible contracts, where either the origin or the destination is left unspecified. Since the vast majority of contracted trade in my sample period is accounted for by the traditional contracts with a fixed origin-destination, I do not consider flexible long-term contracts in my analysis. Table 1 further illustrates characteristics of both the fixed origin-destination and flexible long-term contracts, with the former further sub-divided into ex-ante contracts (signed before the FID date) and ex-post contracts (signed after the FID date). Ex-ante contracts tend to be longer than ex-post contracts (by about 5 years) and specify larger annual quantities to be traded (by around 40%); as such, the total quantity to be traded during the lifetime of a contract is almost 90% higher for ex-ante than ex-post contracts. Ex-ante contracts also require the parties to commit to trade much further in advance: on average, ex-ante contracts are signed 5.3 years before deliveries begin, compared to 2.8 years for ex-post contracts.

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

3 Descriptive Evidence

This section provides descriptive evidence to highlight the role of bargaining power in driving contracting decisions, as well as the potential for short-run misalloacation from contractual rigidities.

Contract timing, investment and bargaining power: Figure 3 plots a histogram of the gap between the contract signature date and the final investment decision (FID) date. A large share of contracts are signed prior to the FID date. Moreover contract decisions are clustered in the period leading up to the investment decision, whereas not as many contracts are signed after the investment decision has been finalized. This suggests sellers prefer to sign contracts with buyers before committing to sunk investments. Figure 4 shows that on average, a seller signs ex-ante contracts amounting to slightly over 60% of their capacity before they commit to the investment, whereas ex-post contracts only account for 23% of capacity.

Figure 3: Gap between contracting date and investment date

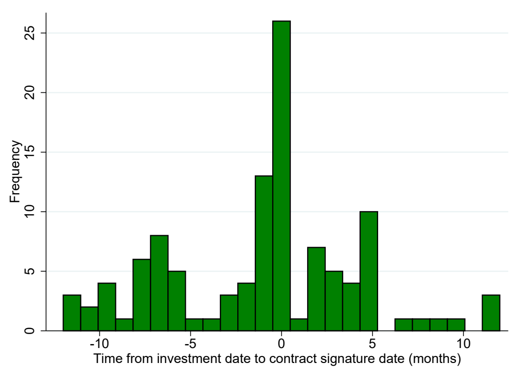

Figure 4: Cumulative share of capacity signed under long-term contracts

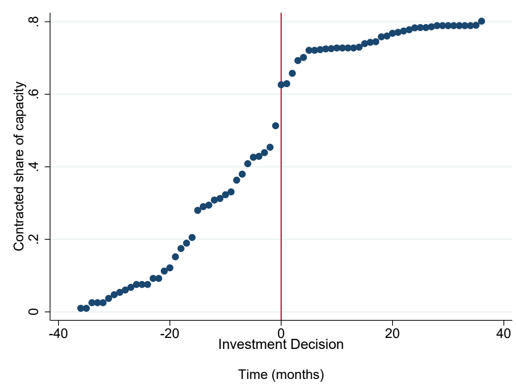

Note: Figure 3 shows the number of months between the contract signature date and the date at which sellers make a final investment decision. Negative values indicate that the contract was signed prior to investment; positive values indicate the contract was signed after the investment decision was made. Figure 4 plots the average share of capacity signed under long-term contracts, against the number of months relative to the date of the final investment decision.

Why is there such widespread use of ex-ante long-term contracts, despite the fact that they require both buyers and sellers to commit to trade several years in advance? A key distinction between ex-ante and ex-post contracting is that by the time ex-post contracts are negotiated, the cost of investment is sunk and cannot directly influence the negotiated price. If the seller is in a weaker bargaining position relative to the buyer, the price they are able to secure in ex-post contract negotiations may not be sufficient to induce them to choose the efficient level of investment, resulting in under-investment and creating incentives to contract ex-ante.

For such an explanation to hold, LNG buyers must possess some degree of bargaining power in contract negotiations; if sellers can costlessly switch between buyers, they would have no reason to fear that any one buyer can capture a significant share of the surplus. Buyer bargaining power is plausible as there is only a limited pool of buyers to contract with at any given point in time (with an average of 12 buyers negotiating long-term LNG contracts in an average year). There is also widespread heterogeneity in contract sizes signed within the same year. As such, switching to a different buyer is likely to entail a large adjustment in the contract quantity, making it costly for the seller to switch to an alternative buyer, as discussed further in Appendix A.2.

The bargaining power of sellers relative to buyers depends significantly on geography, given the importance of shipping costs (which are on average 15% of the final price of LNG). If a seller and buyer are located close to one another, but the seller is located far from other potential alternative buyers, then the seller has to incur higher shipping costs in order to trade with a different buyer (whether via a long-term contract or on the spot market), weakening the seller's bargaining power in contract negotiations. Conversely, when the buyer is located far away from potential alternative sellers, then the buyer's outside option is worse, strengthening the seller's bargaining position. The theoretical model developed later predicts that in such a situation, sellers and buyers should sign larger ex-ante contracts, as a way to mitigate under-investment.

Next, I test if geography affects contracting behavior in the way suggested by the theory. To measure how geography shapes outside options, I first compute how far away each agent signing a long-term contract is from their alternative or next-best trading partners. I use the 25th percentile of the distance between that agent and all potential trading partners as a measure of how far the agent is from alternative trading partners.[14] I then compute the *relative distance* of each agent to their alternative trading partners: this is defined as the distance from the agent to its alternative trading partners (as defined above), divided by the distance to the trading partner that they are negotiating the contract with. The larger the relative distance of the agent to their alternative trading partners, the larger the increase in shipping costs they have to incur in order to switch to a different trading partner, and therefore the weaker their potential bargaining leverage.

I then regress the logarithm of the total contract quantity agreed between sellers and buyers on the relative distance of the seller to their alternative buyers (which measures the strength of the seller's outside option) and relative distance of the buyer to their alternative seller (which measures the strength of the buyer's outside option). I allow these coefficients to differ for ex-ante and ex-post contracts. As Table 2 (Spec. 2) shows, consistent with the theory, sellers and buyers tend to sign *larger* ex-ante contracts if the relative distance of the seller is high (and therefore the seller's outside option is weaker). Likewise, sellers and buyers sign *smaller* ex-ante contracts if the relative distance of the buyer is high (and therefore the buyer's outside option is weaker, giving the seller more bargaining leverage). These effects are sizeable in magnitude: raising the relative distance of

[14]The idea behind using the 25th percentile is that, in the event of a contractual impasse, agents are likely to trade with other trading partners located relatively close to them, but not necessarily the trading partner closest to them (since they may not necessarily be available to trade). The results however are similar if we use instead the median or mean distance from the agent to other potential trading partners.

Table 2: Contract quantity regressions

	Spec. 1		Spec. 2	
	Estimate	S.E.	Estimate	S.E.
Ex-ante contract	0.70	(0.14)	0.64	(0.26)
Distance	0.082	(0.036)	0.072	(0.036)
Relative distance, seller	0.28	(0.11)	0.25	(0.12)
Relative distance, buyer	-0.037	(0.062)	-0.0058	(0.063)
Ex-ante*Relative distance, seller			0.92	(0.32)
Ex-ante*Relative distance, buyer			-0.96	(0.26)
Extension	-0.30	(0.16)	-0.29	(0.16)
Time Trend	-0.022	(0.0055)	-0.024	(0.0055)
N	337		337	
R^2	0.15		0.19	

Note: Each observation is a long-term contract. The dependent variable is the logarithm of the total contract quantity to be traded between the buyer and the seller over the lifetime of the contract. The sample includes every long-term contract that specifies a fixed export and import location (contracts with "flexible" origins or destinations are excluded). Extensions are renewals of existing contracts.

the seller from alternative buyers by 10% raises the average size of ex-ante contracts by 9.2%, while reducing the relative distance of the buyer from alternative sellers by 10% reduces the average size of ex-ante contracts by 9.6%. Appendix A shows that we obtain similar results with alternative definitions of relative distance, or if we control for other determinants of contract quantity (such as export capacity, rule of law indicators and past contracting).

These results indicate that a key role of ex-ante contracting is to allow sellers and buyers to avoid the under-investment that would arise if sellers relied exclusively on ex-post contracts or spot sales. This does not, of course, mean that other motives for long-term contracting (such as supply assurance motives) are not important. The structural model developed later allows me to quantify the extent to which bargaining power and the risk of under-investment drives contracting behavior, while accounting for these other motives.

Long-term contracts, market power and short-run allocation: Next I present evidence suggesting both that there is misallocation in the LNG industry, and that market power and contract rigidities contribute to this. Figure 5 shows spot prices in various LNG importing regions. A striking feature of the industry is the presence of large and systematic spot price differentials across regions, especially during periods when the LNG market is tight. This is exemplified by 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, converging again only in late 2014. Similarly, 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.

Figure 5: LNG Spot Prices in Different Regions

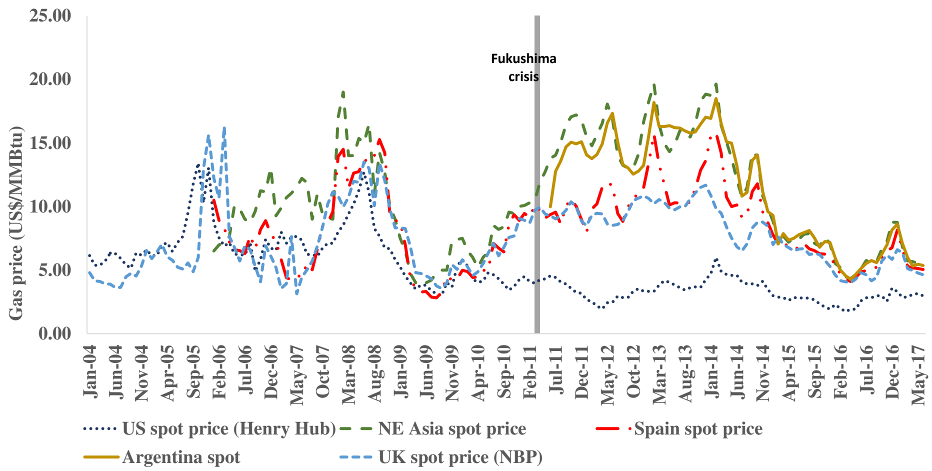

Note: The figure plots monthly spot LNG and gas prices in different regions. The US price is the price of natural gas traded on the Henry Hub. The UK price is the price of natural gas traded on the NBP Virtual Trading Point. The price in North-east Asia is a benchmark spot price for the region (comprising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

In a competitive market with no distortions, faced with these divergent prices, capacity-constrained sellers would sell only to the destination with the highest price net of transportation costs. This would mean that during these periods of high Asia-Europe spot price differentials, spot exports should be mostly directed to Asia, since even after accounting for transportation costs most sellers received higher prices from selling to Asia than to Europe. But as Appendix Figure A2b shows, during the “boom" period of 2011-14, Europe continued to import a significant amount of LNG on the spot market, often 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% from Egypt and Algeria, all countries with roughly similar shipping costs of selling LNG to Asia and Europe. This is inconsistent with competitive behavior by LNG spot sellers. During this period, Europe also continued to import large amounts of contracted LNG, as shown by Appendix Figure A2a, despite having a lower willingness-to-pay than Asian buyers (as indicated by the spot price differentials). This suggests that rigid long-term contracts may also have impeded the market’s response to the demand shock.

These data patterns suggest that the allocation of LNG is likely inefficient. The empirical model developed in the paper quantifies the size of this inefficiency and the extent to which it is explained by contractual rigidities and by seller market power.

4 Model

In this section, I describe a model of contracting, investment and spot trade in the LNG industry. I begin by providing an overview of the model in Section 4.1, where I also introduce the notation and describe the timing of the game. Section 4.2 develops a model of the spot market, while the model of contracting and investment is described in Section 4.3. Finally, Section 4.4 discusses the key equilibrium properties of the model as well as the implications of the main modelling assumptions.

4.1 Model overview and timing

Time is discrete and indexed by $t = 1, ..., T$. Each period denotes a year. There are J buyers indexed by $j = 1, ..., J$, and N sellers indexed by $i = 1, ..., N$. Sellers produce a homogeneous good (LNG). Buyers and sellers are risk-neutral and have discount factor β.

In the empirical analysis, each exporting country is treated as a separate seller, and each importing country is treated as a separate buyer. Buyers and sellers are spatially differentiated, and d_{ij} is the distance between seller i and buyer j. Buyers have uncertain demand for LNG.

Each seller i owns an export project, and decides how much capacity K_i to build. Once the construction of the project is complete, the capacity becomes available on the market and the seller can begin producing and exporting LNG, both via long-term contracts and on the spot market.[15]

I model contracting and investment decisions for each seller i (and associated buyers) as a sequential, multi-stage game, as summarized in Figure 6.

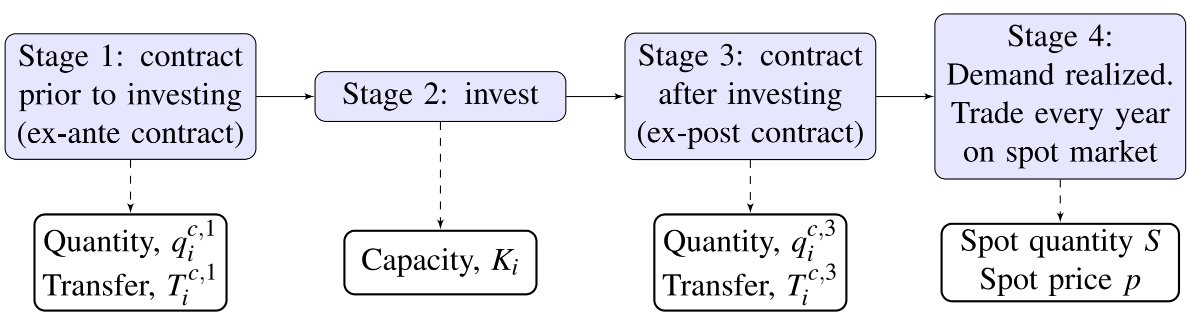

Figure 6: Stages of the game

Stage 1 of the game takes place before the seller has made any investment decision. Each seller-buyer pair bargains over the contract quantity and a lump-sum transfer to be paid by the buyer to the seller. In Stage 2, the seller chooses how much to invest, taking as given any contracts already signed in Stage 1. The seller must build at least enough capacity to meet these pre-committed quantities. Stage 3 of the game occurs after the seller has committed to the investment. Just as in

[15]For notational simplicity, I describe the model for the case where each seller makes only a single investment. In the more general version of the model, which I use for estimation, a seller may make multiple, distinct investments.

Stage 1, each seller-buyer pair then bargains over the contract quantity and a lump-sum transfer.[16] Finally, every year, all sellers and buyers participate in the LNG spot market (Stage 4). Demand shocks are realized and sellers meet their contractual obligations. Sellers can sell any uncommitted capacity on the annual spot market.

The model features the key economic mechanisms underlying contractual choices. Firstly, because the seller bargains with the buyer over how to split the surplus from contracting, *ex-ante* and *ex-post* contracting have differing implications for sellers' investment incentives. If sellers only rely on ex-post contracting (Stage 3), then they cannot recoup the full marginal value of their investment as long as buyers have some bargaining power, and will under-invest. Signing contracts ex-ante (in Stage 1) allows the seller and buyer to forestall under-investment, as suggested by the theoretical literature on contracting (Williamson, 1975; Edlin and Reichelstein, 1996). However, both the magnitude of potential under-investment and the resulting incentive to sign ex-ante contracts depend critically on the seller and buyers' access to other trading partners via the spot market, with under-investment less of a concern for sellers with especially strong outside options.

Secondly, the use of long-run contracts has implications for allocative efficiency. Because contract quantities are fixed and have to be agreed before demand shocks are realized, long-term contracts reduce the flexibility of the market in responding to demand shocks, which can result in short-run misallocation. At the same time, because sellers exercise market power on the spot market, long-term contracts also have pro-competitive effects, as described by Allaz and Vila (1993). I return to these and other properties of the model in Section 4.4. In the next two sections, I describe the model in reverse order of timing, starting with Stage 4.

4.2 Demand, production and spot trade

This section describes a model of LNG spot trade. By this stage of the game, all investment and contracting decisions have been made. Buyer demand shocks are now realized. Sellers observe these shocks and choose how much to sell to different buyers on the spot market.

Demand: The demand for buyer j at time t is given by the following equation:

$$Q_{jt} = Q_d(p_{jt}, R_{jt}, x_{jt}, \varepsilon_{jt}) \tag{1}$$

where p_{jt} denotes the spot price paid by buyer j, x_{jt} are demand shifters such as weather and the price of competing fuels, and ε_{jt} denotes demand shocks. R_{jt} is a measure of buyer capacity which I describe more fully in Section 5. Q_{jt} is the *total* quantity of LNG purchased by buyer j, including both long-term contracts and purchases on the spot market.

Note that each buyer faces a different price p_{jt}, since buyers are spatially differentiated. So in

[16]These ex-post contracts may be negotiated either before or after the seller has completed construction of the project

effect there are J different spot markets, each with its own spot price, though these prices will be correlated since the same sellers can sell to each market.

Trade Flows: The quantity of LNG sold by seller i to buyer j equals $q_{ijt} = S_{ijt} + q^c_{ijt}$ where S_{ijt} denotes spot sales and q^c_{ijt} denotes contracted sales. Seller i's total production equals $q_{it} = \sum_j q_{ijt} = \sum_j (S_{ijt} + q^c_{ijt})$, while buyer j's total imports equal $Q_{jt} = \sum_i q_{ijt} = \sum_i (S_{ijt} + q^c_{ijt})$.

Costs of production and sales: Sellers incur costs both in producing LNG, and then in shipping the LNG to the buyers. Let $C(q_{it}, K_{it})$ denote the convex cost of production, where K_{it} is seller i's total capacity in period t. Production cost depends on K_{it} since capacity constraints are significant in this industry (as evidenced by Appendix Figure A1) and sellers with larger capacity are able to export more LNG, indicating that capacity is an important determinant of marginal costs. In addition, the seller incurs shipping costs that differ by buyer. Each unit of LNG costs c^d_{ijt} to ship from i to j, where the shipping cost c^d_{ijt} is increasing in the distance d_{ij} between seller i and buyer j and fluctuates by year.[17] Thus firm i's total costs from LNG sales of $\{q_{ijt}\}_{j=1}^J$ are given by:

$$C_i(\{q_{ijt}\}_{j=1}, K_{it}) = C(q_{it}, K_{it}) + \sum_j c^d_{ijt} q_{ijt}$$

Spot market equilibrium: All J markets clear separately and simultaneously. The market clearing price vector $p^*_t = (p^*_{1t}, ..., p^*_{jt}, ..., p^*_{Jt})$ is characterized by the following set of equations:

$$Q_d(p^*_{jt}, R_{jt}, x_{jt}, \varepsilon_{jt}) = \sum_{i=1}^N q^c_{ijt} + \sum_{i=1}^N S_{ijt} \ , \forall j$$

Sellers engage in Cournot competition. In a Cournot equilibrium, each seller i takes as given rival spot quantities $\{S_{-ijt}\}_{j=1}^J$ and chooses the vector of spot quantities, $\{S_{ijt}\}_{j=1}^J$, that maximizes its sum of profits across all markets:[18]

$$\{S_{ijt}\}_{j=1}^J = argmax_{\{\hat{S}_{ijt}\}_{j=1}^J} \left[\sum_{j=1}^J p^*_{jt}(S_{ijt}, S_{-ijt}) \hat{S}_{ijt} - C(q_{it}, K_{it}) - \sum_j c^d_{ijt} q_{ijt} \right]$$

The first-order condition satisfied by the optimal quantity S_{ijt} (with equality if $S_{ijt} > 0$) is:

$$\underbrace{p^*_{jt} + S_{ijt} \frac{\partial p^*_{jt}(S_{ijt}, S_{-ijt})}{\partial S_{ijt}}}_{\text{Marginal revenue of selling to market } j} - \underbrace{\left(\frac{\partial C(q_{it}, K_{it})}{\partial S_{ijt}} + c^d_{ijt} \right)}_{\text{Marginal cost of selling to market } j} \leq 0 \qquad (2)$$

[17] c^d_{ijt}, which is directly observed in the dataset, depends on d_{ij} as well as prevailing LNG shipping rates.

[18] The seller's payoff function does not include revenue from contracted sales, since this is unaffected by spot market decisions and is effectively "sunk" at this time.

Per-period payoffs: The payoffs in period t are functions of q_t^c (a vector of all contracted trade flows), K_t (a vector of capacities) and ε_t (a vector of demand shocks). Integrating out the demand shocks, we can derive expected per-period payoffs to sellers and buyers as functions of q_t^c and K_t. Let $\pi_{it}^s(q_t^c, K_t)$ denote seller i's expected payoff (variable profit) in period t. Similarly, let $\pi_{jt}^b(q_t^c, K_t)$ denote buyer j's expected payoff (consumer surplus) in period t.

Discussion: Because of capacity constraints, the seller's choice of S_{ijt} (i.e., their spot sales in market j) depends on how much they sell in other markets. If the seller sells high quantities in some other market (e.g., due to a demand shock), then the seller's marginal cost of production increases and so the seller will sell less in market j. This in turn implies that demand shocks in one market will affect firms' spot market sales across all markets. The *interconnectedness* of markets is an important feature of the model: if instead sellers faced constant returns to scale, their decisions in one market would be independent of their decisions in other markets.

I assume sellers are strategic and engage in spatial price discrimination on the spot market.In a companion paper, I test whether this is true, finding that seller behavior is consistent with Cournot competition and unlikely to be generated by perfect competition (Zahur, 2022).

I assume that buyers cannot engage in arbitrage on the spot market: spatial price discrimination would not be feasible if costless arbitrage were possible. But there are a number of barriers to arbitrage in LNG. The most important of these are resale restrictions that are commonly written into long-term LNG contracts, such as destination clauses and diversion clauses. Physical arbitrage (i.e., re-export of LNG) is costly because of shipping costs and additional costs incurred in re-exports (such as the cost of reloading LNG from one tanker to another). Finally, financial markets for LNG have been historically very limited, and derivatives trade is small compared to physical trade, further limiting the scope for arbitrage: this is discussed further in Appendix A.1.

4.3 Contracting and Investment

I now embed this model of short-run LNG flows into an equilibrium model of long-run contracting and investment (Stages 1 - 3 of the multi-stage game).

Setup: Let $\mathbf{B}_i^1$ denote the exogenously determined set of buyers with whom seller i can sign ex-ante contracts (in Stage 1 of the game). Let $\mathbf{B}_i^3$ denote the set of buyers with whom the seller can sign ex-post contracts (in Stage 3 of the game). These sets of buyers may overlap.

A long-term contract signed between seller i and buyer j consists of a start date, an end date, a contracted quantity q_{ij}^c (to be delivered by the seller to the buyer for every year during which the contract is operational) and a lump-sum transfer T_{ij}^c which the buyer pays to the seller. All contracts are assumed to include destination clauses preventing re-sales, and the contract quantity cannot be re-negotiated from one year to one another. Thus in a given year t, the total contracted quantity seller i delivers to buyer j, q_{ijt}^c, is the sum of contracted quantities across all contracts signed between seller i and buyer j that are still active in period t.

Let $q_i^{c,1} = \left\{q_{ij}^c\right\}_{j\in\mathbf{B}_i^1}$ and $q_i^{c,3} = \left\{q_{ij}^c\right\}_{j\in\mathbf{B}_i^3}$ denote vectors comprising all the contract quantities signed by seller i in Stage 1 and Stage 3 respectively.

Beliefs and expected payoffs: A key assumption I make is that shocks to demand, ε_{jt}, are not serially correlated, so that the realization of ε_j in period t is independent of buyer j's past realizations of ε_j. This implies that agents do not learn about demand with time. As such, agents can perfectly foresee the future contracting and investment decisions of rival agents. While the assumption of no serial correlation in demand shocks is strong, it makes the model tractable to estimate since I do not need to estimate agents' beliefs about the actions of rivals in future periods. Though I assume away serial correlation, I do allow the volatility of demand shocks to be different across countries: specifically, I assume $\varepsilon_{jt} \sim N(0,\sigma_j^2)$, where σ_j is different for each country j; thus, the demand shocks are potentially heteroskedastic.

Under the assumption of no serial correlation, we can now work out the expected payoffs to a particular seller i and the buyers that seller i contracts with. Let $Y_{-i} = (q_{-i}^c, K_{-i})$ be a vector of the contracting and investment decisions made in projects operated by sellers other than i. Seller i and its buyers take Y_{-i} as given when making their own contracting and investment decisions.

Let t_i^3 denote the time period when Stage 3 contracts have been finalized. The expected payoffs to the buyer and seller are their discounted sum of lifetime expected payoffs from period t_i^3 onwards, as functions of their own contract quantities $\left\{q_i^{c,1}, q_i^{c,3}\right\}$ and capacity K_i, as well as contract capacities and capacities chosen by other sellers and buyers (which are summarized in Y_{-i}).

Seller i's expected lifetime payoff, V_i^3, can be written as the discounted sum of their expected profits in each future period:[19]

$$V_i^3(q_i^{c,1}, q_i^{c,3}, K_i, Y_{-i}) = \sum_{t=t_i^3}^{\infty} \beta^{t-t_i^3} \underbrace{\pi_{it}^s(q_t^c, K_t)}_{\text{Expected per-period variable profit}} \tag{3}$$

The buyer's expected payoff includes the discounted sum of their expected consumer surplus in each future period, $\pi_{jt}^b(q_t^c, K_t)$. In addition, I allow for the possibility that buyers may be willing to pay a "premium" to purchase contracted LNG rather than spot LNG. The contract premium is the additional utility buyers get from purchasing LNG using long-term contracts as opposed to the spot market.[20] There are several reasons why a contract premium may exist. Buyers may have supply assurance or security motives for preferring contracts which allow them to lock in a portion of their purchases (Bolton and Whinston, 1993): supply security is frequently cited by LNG buyers as a motive for signing long-term contracts (IEA, 2013). By signing a long-term contract, buyers can avoid ex-ante transaction costs incurred in repeatedly participating on the spot market (MacKay,

[19] The superscript "3" is to indicate that these are their payoffs after Stage 3 is complete.

[20] The premium may be negative for buyers who have a higher willingness-to-pay for spot LNG than contracted LNG.

2022). Long-term contracts also reduce the cost of trading frictions (Tolvanen et al., 2022). Finally, risk-averse buyers may prefer contracts to reduce the volatility of their costs of purchasing LNG.

Let $\omega_j^3(q_{ij}^{c,3}, \eta_{ij}^3)$ denote the contract premium the buyer receives from signing a contract of quantity $q_{ij}^{c,3}$ in Stage 3 with seller i, where η_{ij}^3 is a publicly observable shock to the marginal value of contracting between seller i and buyer j (but unobservable to the econometrician). Then, buyer j's expected lifetime payoff at the end of Stage 3, W_j^3, can be written as the sum of their lifetime expected consumer surplus (which I term $\tilde{W}_j^3$) and their contract premium ω_j^3:

$$
\begin{aligned}
W_j^3(q_i^{c,1}, q_i^{c,3}, K_i, Y_{-i}, \eta_{ij}^3) &= \sum_{t=t_i^3}^{\infty} \beta^{t-t_i^3} \underbrace{\pi_{jt}^b(q_t^c, K_t)}_{\text{Per-period consumer surplus}} + \underbrace{\omega_j^3(q_{ij}^{c,3}, \eta_{ij}^3)}_{\text{Contract premium, Stage 3}} \\
&= \underbrace{\tilde{W}_j^3(q_i^{c,1}, q_i^{c,3}, K_i, Y_{-i}, \eta_{ij}^3)}_{\text{Lifetime consumer surplus}} + \underbrace{\omega_j^3(q_{ij}^{c,3}, \eta_{ij}^3)}_{\text{Contract premium, Stage 3}}
\end{aligned}
\tag{4}
$$

Bargaining model A contract consists of a contract quantity and a lump-sum transfer. At each of the two contracting stages (Stages 1 and 3 in Figure 6), contracts are negotiated via Nash-in-Nash bargaining. Each seller-buyer pair chooses the contract that maximizes the Nash product of their surplus from contracting, assuming that every other pair reaches agreement. Since lump-sum transfers are possible, this implies that each seller-buyer pair will (optimally) choose the contract quantity that maximizes their joint surplus, taking as given the contract quantity chosen by other seller-buyer pairs. They will then negotiate the lump-sum transfer to divide the surplus from trading between the two parties (Chipty and Snyder, 1999).

Stage 3: contracting after investment: I now describe how decisions are made at each of the three stages of the game for each project, starting from Stage 3, where the seller (who has already committed to building a capacity of K_i) negotiates contracts with a set of buyers $\mathbf{B}_i^3$.

Equilibrium Quantities, Stage 3

Each seller-buyer pair in Stage 3 of the game chooses the contract quantity that maximizes their joint surplus, taking as given the choices of other pairs. The equilibrium quantities are given by:

$$
q_{ij}^{c,3} = argmax_{q_{ij}} \left[V_i^3(q_i^{c,1}, q_i^{c,3}, K_i, Y_{-i}) + W_j^3(q_i^{c,1}, q_i^{c,3}, K_i, Y_{-i}, \eta_{ij}^3) \right], \forall j \in \mathbf{B}_i^3
$$

The FOC of the quantity problem is:

$$
\frac{\partial V_i^3}{\partial q_{ij}^{c,3}} + \frac{\partial W_j^3}{\partial q_{ij}^{c,3}} = 0 \tag{5}
$$

Equilibrium transfers, Stage 3

Each seller-buyer pair then chooses a transfer paid by the buyer to the seller to maximize the

Nash product of the seller's surplus and the buyer's surplus, taking as given that all other pairs reach agreement. Following the Horn and Wolinsky (1988) notion of a Nash equilibrium among the Nash bargains, I assume that the contracts are binding even in the event that one or more of the negotiations fail. If negotiations fail between seller i and buyer j, they are unable to negotiate any new contracts to replace the contract they failed to sign. Instead, they move on to the spot market, where they can potentially find other trading partners. Thus the disagreement payoffs for both the seller and buyer are determined by their value from participating in the spot market.

Let $V_i^3(q_i^{c,1}, q_{i,\setminus ij}^{c,3}, K_i, Y_{-i})$ denote seller i's disagreement payoff when negotiating with buyer j, where $q_{i,\setminus ij}^{c,3}$ denotes the vector of contract quantities if we set $q_{ij}^{c,3}$ to 0 but maintain all other contracts in $q_i^{c,3}$. Similarly $W_j^3(q_i^{c,1}, q_{i,\setminus ij}^{c,3}, K_i, Y_{-i})$ denotes buyer j's disagreement payoff when negotiating with seller i. Then the lump-sum transfers $T_{ij}^{c,3}$ that seller i receives from buyer j in Stage 3 are given by:

$$T_{ij}^{c,3} = argmax_T \big(\underbrace{V_i^3(q_i^{c,1}, q_i^{c,3}, K_i, Y_{-i}) - V_i^3(q_i^{c,1}, q_{i,\setminus ij}^{c,3}, K_i, Y_{-i}) + T}_{\text{Seller's gains from trade in Stage 3}}\big)^{\tau}$$
$$\big(\underbrace{W_j^3(q_i^{c,1}, q_i^{c,3}, K_i, Y_{-i}, \eta_{ij}^3) - W_j^3(q_i^{c,1}, q_{i,\setminus ij}^{c,3}, K_i, Y_{-i}) - T}_{\text{Buyer's gains from trade in Stage 3}}\big)^{1-\tau}, \forall j \in \mathbf{B}_i^3 \tag{6}$$

where τ is the bargaining weight of seller i when negotiating with buyer j.[21] The higher the seller's bargaining power τ, the larger the transfer that the buyer pays to the seller.

Expected payoffs at the end of Stage 2

The expected payoff for seller i at the end of Stage 2, V_i^2, equals the sum of their expected payoff at the end of Stage 3, V_i^3 and transfers received from the buyers with whom the seller contracts in Stage 3. The payoffs for any buyer j that contracts in Stage 3, W_j^2, is the expected payoff at the end of Stage 3, W_j^3, minus any transfers made to the seller.

$$V_i^2(q_i^{c,1}, K_i, Y_{-i}, \eta_{ij}^3) = V_i^3(q_i^{c,1}, q_i^{c,3}, K_i, Y_{-i}) + \sum_{j \in \mathbf{B}_i^3} T_{ij}^{c,3}(q_i^{c,1}, q_i^{c,3}, K_i, Y_{-i}, \eta_{ij}^3)$$
$$W_j^2(q_i^{c,1}, K_i, Y_{-i}, \eta_{ij}^3) = W_j^3(q_i^{c,1}, q_i^{c,3}, K_i, Y_{-i}, \eta_{ij}^3) - T_{ij}^{c,3}(q_i^{c,1}, q_i^{c,3}, K_i, Y_{-i}, \eta_{ij}^3)$$

Stage 2: investment In Stage 2, the seller chooses how much capacity to build. Let $\Gamma_i(K_i, \eta_i^2)$ denote the cost of the investment. η_i^2 is a publicly observable shock to the marginal cost of investing by seller i. The seller chooses K_i to maximize their net lifetime benefit from investing:

$$K_i^* = argmax_{K_i}\big[V_i^2(q_i^{c,1}, K_i, Y_{-i}, \eta_{ij}^3) - \Gamma_i(K_i, \eta_i^2)\big] \tag{7}$$

[21] In robustness checks, I allow the bargaining weight to differ by seller and buyer.

The first-order condition to the seller's investment problem is:

$$\frac{\partial V_i^2}{\partial K_i} - \frac{\partial \Gamma_i(K_i, \eta_i^2)}{\partial K_i} - = 0 \tag{8}$$

Stage 1: contracting prior to investment In Stage 1, just as with Stage 3, I allow for the buyer to receive a premium from signing a contract of quantity $q_{ij}^{c,1}$ with seller i, to capture differences in buyer preferences between contracted and spot LNG. The size of this contract premium in Stage 1 may differ from that in Stage 3. The contract premium term equals $\omega_j^1(q_{ij}^{c,1}, \eta_{ij}^1)$, where η_{ij}^1 is a publicly observable shock that affects the marginal value from contracting between seller i and buyer j.

Let V_i^1 denote seller i's lifetime expected payoff for any set of contract quantities in Stage 1. Let W_j^1 denote buyer j's lifetime expected payoff from a vector of contract quantities in Stage 1. These payoffs equal:

$$V_i^1(q_i^{c,1}, Y_{-i}, \eta_i^2, \eta_{ij}^3) = V_i^2(q_i^{c,1}, K_i, Y_{-i}, \eta_{ij}^3) - \Gamma_i(K_i, \eta_i^2)$$
$$W_j^1(q_i^{c,1}, Y_{-i}, \eta_{ij}^1, \eta_i^2, \eta_{ij}^3) = W_j^2(q_i^{c,1}, K_i, Y_{-i}, \eta_{ij}^3) + \omega_j^1(q_{ij}^{c,1}, \eta_{ij}^1)$$

Equilibrium Quantities, Stage 1

Each seller-buyer pair chooses the contract quantity that maximizes their joint payoff in a Nash equilibrium played with the other pairs. The equilibrium quantities are:

$$q_{ij}^{c,1} = argmax_{q_{ij}} \left[V_i^1(q_i^{c,1}, Y_{-i}, \eta_i^2, \eta_{ij}^3) + W_j^1(q_i^{c,1}, Y_{-i}, \eta_{ij}^1, \eta_i^2, \eta_{ij}^3) \right], \forall j \in \mathbf{B}_i^1$$

The FOC is:

$$\frac{\partial V_i^1}{\partial q_{ij}^{c,1}} + \frac{\partial W_j^1}{\partial q_{ij}^{c,1}} = 0 \tag{9}$$

Equilibrium Transfers, Stage 1

I assume that if bargaining between i and j breaks down, the agents move on to the next stage of the game. Let $V_i^1(q_{i,\setminus ij}^{c,1}, Y_{-i}, \eta_i^2, \eta_{ij}^3)$ denotes seller i's disagreement payoff when negotiating with buyer j, where $q_{i,\setminus ij}^{c,1}$ denotes the vector of Stage 1 contract quantities if we set $q_{ij}^{c,1}$ to 0 but maintain all other contracts in $q_i^{c,1}$. Similarly $W_j^1(q_{i,\setminus ij}^{c,1}, Y_{-i}, \eta_i^2, \eta_{ij}^3)$ denotes buyer j's disagreement payoff when negotiating with seller i.

The lump-sum transfers T_{ij}^c that seller i receives from buyer j in Stage 1 are determined by

Nash-in-Nash bargaining, with the equilibrium transfers characterized by the following equation:

$$T_{ij}^{c,1} = argmax_T \big(\underbrace{V_i^1(q_i^{c,1}, Y_{-i}, \eta_i^2, \eta_{ij}^3) - V_i^1(q_{i,\setminus ij}^{c,1}, Y_{-i}, \eta_i^2, \eta_{ij}^3) + T}_{\text{Seller's gains from trade in Stage 1}}\big)^{\tau}$$

$$\big(\underbrace{W_j^1(q_i^{c,1}, Y_{-i}, \eta_{ij}^1, \eta_i^2, \eta_{ij}^3) - W_j^1(q_{i,\setminus ij}^{c,1}, Y_{-i}, \eta_i^2, \eta_{ij}^3) - T}_{\text{Buyer's gains from trade in Stage 1}}\big)^{1-\tau}, \forall j \in \mathbf{B}_i^1$$

4.4 Equilibrium properties of the contracting and investment game

Bargaining power, contracting and investment: A key property of the model is that the seller's incentive to invest depends on their bargaining power relative to the buyer. If the seller has limited bargaining power and can only sign contracts ex-post, they will under-invest, since they do not fully internalize the benefits realized by the buyer from their investment.

The magnitude of this hold-up effect depends on the seller's Nash bargaining weight τ: the smaller τ is, the smaller the share of surplus captured by the seller, and the more severe the under-investment. The extent of under-investment also depends on the relative strength of the outside options of the seller and buyer, with the two having opposite effects on investment. As the seller's outside option weakens, or the buyer's outside option strengthens, the price negotiated in ex-post contracts decreases, worsening under-investment. A consequence of this is that seller market power in the spot market can be socially beneficial, by reducing the incentives of sellers to under-invest. When sellers have more market power, their outside options are stronger and buyer's outside options are weaker; both of these increase the bargaining leverage of sellers in contract negotiations, leading to reduced under-investment. The presence of these outside options does not entirely eliminate under-investment, however. The seller's investment is efficient only in the polar case where the seller has the ability to fully capture the surplus from trade when negotiating with the buyer. Appendix D.1 provides further details as well as results from numerical simulations of a simplified version of the model that illustrates these predictions.

Bargaining power and ex-ante contracting: Foreseeing potential under-investment in Stage 2, the seller and buyer(s) have an incentive to sign ex-ante contracts in Stage 1, as a way to induce the seller to invest more (to the mutual benefit of both parties). Ex-ante contracts differ in a fundamental way from ex-post contracts, since the seller has not yet paid the sunk cost of investment. Thus, when the seller and buyer negotiate an ex-ante contract, they will choose the contract quantity to maximize their joint surplus, which includes the cost of the investment. The use of ex-ante contracts will therefore reduce the extent to which sellers under-invest (Edlin and Reichelstein, 1996). It follows that the more severe the potential risk of under-investment, the larger the size of ex-ante contracts signed between the buyer and seller. Numerical simulations of the model reported in Appendix D.1 confirm this intuition: I find that as the potential under-investment becomes more severe (either because the seller has a lower bargaining weight τ or because the seller's outside option is weak relative to the buyer), firms sign larger ex-ante contracts and smaller ex-post contracts.

The effectiveness of using ex-ante contracts to guard against under-investment in this fashion depends on whether there are any additional costs from contracting ex-ante. If the value of the contract to the buyer and seller is the same regardless of when it is signed (conditional on a given level of investment), then sellers and buyers will only sign contracts ex-ante, which will completely eliminate under-investment. This will happen, for instance, if the contract premia in Stage 1 and Stage 3 are equal (that is, if $\omega_j^1 = \omega_j^3$). But if there are any additional benefits from contracting ex-post instead of ex-ante (so that $\omega_j^3 > \omega_j^1$), sellers and buyers will use a mix of ex-ante and ex-post contracting, and ex-ante contracting will reduce but not completely eliminate under-investment.

It is worth emphasizing here that mitigating the risk of under-investment is not the only reason to sign long-term contracts. As mentioned earlier, long-term contracts may be negotiated because buyers are willing to pay a premium for contracted purchases than spot purchases, due to supply assurance concerns, to lower transaction costs and due to risk aversion. The empirical analysis that I carry out will seek to disentangle these different motives for contracting.

Allocative efficiency and contracting externalities: The use of long-term contracts also has consequences for short-run allocative efficiency. There are two competing forces at work. On the one hand, long-term contracts require sellers to commit to selling to pre-selected buyers at a time when demand is not fully known. This can reduce the ability of firms to flexibly respond to ex-post demand shocks. The lack of flexibility is exacerbated by the presence of contractual resale restrictions (such as destination clauses) that limit the ability of buyers of contracted LNG to re-sell to other buyers that may be experiencing increased demand. On the other hand, as argued by Allaz and Vila (1993), contracts can have a pro-competitive effect when sellers have market power on the spot market. The intuition behind this is that once firms have negotiated contracts, they act more aggressively in the spot market, since they have less to lose by driving the spot price down (since the price they receive on their contracted portion of their output is unaffected by the spot price).

The trade-off between flexibility and market power depends on how capacity-constrained sellers are and the extent of their contractual commitments. If firms produce well below capacity, contracts signed with one buyer do not hurt their ability to meet demand shocks experienced by other buyers; here, we would expect the Allaz and Vila (1993) market power effect to dominate, so that contracts improve allocative efficiency. But if firms produce at close to full capacity and if contracts account for a large share of their capacity, contracts signed with one buyer can severely limit their ability to meet demand shocks experienced by other buyers. In this case we would expect the flexibility effect to dominate, so that contracts hurt allocative efficiency. Appendix D.2.1 reports results from Monte Carlo simulations of the model that illustrate both these effects.

Contracting Externalities: Long-term contracts also impose externalities on buyers and sellers who are not party to the contract. A rich theoretical literature has shown that contracting externalities are pervasive when bilateral contracts are used (Hart and Tirole, 1990; Bolton and Whinston, 1993; Segal, 1999). The intuition is simple: the two parties that contract maximize their bilateral surplus from trade, but do not internalize the impact (whether positive or negative) on other agents.

In this particular context, contracting externalities arise in a few ways. A contract signed between a seller and one buyer can impose negative externalities on “excluded" buyers who are not part of the contract. This is because the contract commits a part of the seller’s output for the exclusive use of one buyer and reduces the quantity of LNG available to the excluded buyers on the spot market, who end up paying higher spot prices (especially during periods of high demand). This negative externality is most salient when sellers are capacity constrained and buyers are competing for limited supplies.[22] However, the negative externality on buyers is less significant if the seller signing the contract is not capacity-constrained, since the contract is then unlikely to affect their marginal cost of supplying to other buyers. Long-term contracts also impose externalities on excluded sellers, since sellers directly compete with each other on the spot market. Appendix D.2.2 presents Monte Carlo simulations of the model that illustrate how contracts unambiguously impose negative externalities on excluded buyers as long as sellers are capacity constrained. Contracts may impose positive or negative externalities on excluded sellers, depending on who the contract is signed with and how the contract affects the level of competition sellers face in their primary markets.

When contracting externalities exist, the equilibrium level of contracting will generally not be socially optimal. If the externalities are primarily negative, there will be over-contracting in equilibrium; if the externalities are primarily positive, there will be under-contracting. Whether or not the level of contracting is socially optimal is therefore an empirical question that I return to after estimating the structural model.

Discussion of modelling assumptions: Like most of the empirical literature on bargaining, I assume Nash-in-Nash bargaining, which has a few well-known limitations. The one that is most salient in this setting is that if a seller and a buyer cannot agree to a contract in Stage 3, they are unable to replace the failed contract by signing new contracts with other buyers and sellers, and must instead trade on the spot market for any remaining quantities they wish to purchase/sell.[23] An alternative approach developed by Ho and Lee (2019), Nash-in-Nash with Threat of Replacement, allows one of the firms to threaten to replace their trading partner with a different trading partner in the event of disagreement. However, their approach permits only one side of the market to exercise the threat of replacement. In the LNG industry, both sellers and buyers may have credible outside

[22]This effect is similar to the negative externality identified by Bolton and Whinston (1993), who show that vertical integration between a capacity-constrained seller and a buyer harms other, non-integrated buyers.

[23]This is less of an issue with Stage 1 contracting, since the seller and buyer can respond to disagreement in Stage 1 by signing larger contracts with other trading partners in Stage 3.

options: sellers might be able to contract with other buyers, and buyers with other sellers. To my knowledge, tractable empirical models of bargaining that allow both parties to exercise outside options involving replacement have not yet been developed.

Sellers and buyers are assumed to negotiate a lump-sum transfer together with the contract quantity. In practice, they negotiate a pricing formula that indexes the LNG price to the price of a benchmark, usually the oil price (Agerton, 2017). As long as agents are risk-neutral, though, negotiating a pricing formula is equivalent to negotiating a lump-sum transfer equal to the expected discounted sum of payments that the buyer makes under any given pricing formula.

I assume that long-term contracts cannot be re-negotiated or breached. Re-negotiation is rare in the LNG industry: in the dataset, only 7 out of 464 contracts were re-negotiated. Some contracts include "price review" clauses that allow the parties to periodically re-negotiate the pricing formula, but successful price re-negotiations have historically been uncommon in the industry, and the renegotiation process is costly and time-consuming (Ason, 2019). Re-negotiation of the contract quantity or destination flexibility is even more uncommon, with contracts structured to provide very limited ability to adjust these terms (IEA, 2013). Likewise, Weems (2016)'s analysis of LNG disputes found relatively few instances of buyers or sellers breaching long-term contracts. Breaching LNG contracts is costly both because of the negative reputational consequences from breaching and because of the risk of having to pay breach remedies.[24] The use of indexed (rather than fixed) prices also limits incentives to breach contracts (Blouin and Macchiavello, 2019).

The fact that LNG contracts cannot easily be breached or re-negotiated ex-post implies that ex-ante contracting is effective as an instrument for reducing under-investment. By contrast, in markets where contract enforcement is weak, there is a significant risk that parties will breach contracts and/or opportunistically renegotiate contracts ex-post. In such settings, under-investment can be significant even when ex-ante long-term contracts are used (Ryan, 2020, 2021).

In the model developed in this paper, firms guard against under-investment by signing ex-ante long-term contracts. An alternative organizational remedy is vertical integration (Grossman and Hart, 1986). If sellers and buyers could vertically integrate at no cost, that would fully eliminate any under-investment. Full vertical integration, however, is rarely observed in the LNG industry, as discussed further in Appendix A. Vertical integration is legally infeasible in many countries that require their LNG export projects to be majority owned by a domestic firm (often the national oil company). Furthermore, long-term contracts allow a seller to sell to multiple buyers, which may be difficult to achieve via vertical integration. As such, I do not model vertical integration.

Another potential benefit of ex-ante contracts is that they may allow financially constrained sellers to lower costs of borrowing. Hartley (2015) emphasizes the role of long-term LNG contracts

[24] For example, in *Statoil v Sonatrach (2013)*, a US$536 million award was issued against Sonatrach by the International Chamber of Commerce (ICC) after it was found to be in breach of their long-term supply agreements with Statoil.

in reducing cash flow variability and thereby increasing the debt capacity of investment projects. While I do not explicitly model this channel, lower financing costs from ex-ante contracting are implicitly captured in the ex-ante contract premium ω_1, which reflect any unobservables affecting the preferences of agents for signing contracts ex-ante. Finally, the model focuses on explaining the size distribution of ex-ante and ex-post long-term contracts, and takes their duration as given: an empirical model of optimal contract duration can be found in MacKay (2022).

5 Estimation

I now describe how to estimate the model described in Section 4 and discuss estimation results.

5.1 Estimation of Demand Curve

The demand curve (1) for country j in period t is parameterized as follows:

$$\frac{Q_{jt}}{R_{jt}} = \alpha - bp_{jt} + \theta_j + x_{jt}\theta_{dx} + \varepsilon_{jt} \tag{10}$$

I use quarterly panel data to estimate demand, so t denotes a year-quarter. R_{jt} is a measure of country j's size. Thus Q_{jt}/R_{jt} is country j's quantity scaled by its size, which I assume is a linear function of the price and demand shifters x_{jt}. In contrast to a standard linear demand function, this demand system allows a given change in spot prices to have a bigger impact on demand for large buyers (e.g., Japan) than for small buyers (e.g., Dominican Republic).[25] The specific measure of R_{jt} I use is the maximum regasification capacity built by country j during the sample period, which is a good proxy for size since larger buyers build more import capacity.[26]

In the baseline specification, x_{jt} includes electricity consumption from fossil fuels, the price of oil and the minimum temperature reached during period t. Electricity consumption from fossil fuels is a measure of the residual demand for electricity, after accounting for baseline generation from sources such as nuclear energy.[27] The minimum temperature reached during period t helps account for the fact that LNG demand spikes during colder winters, when there is greater demand for natural gas for heating. The oil price is included since oil and natural gas are substitutes in power generation, so an increase in the oil price might be expected to raise LNG demand. Finally, I include country fixed effects, θ_j, to capture any remaining time-invariant differences in LNG demand across countries (for example the amount of piped natural gas the country has access to).

I use two instruments for the spot price p_{jt}. The first is the average minimum temperature in rival importers during the same period. When rival importers experience colder weather than expected,

[25] The ideal approach for dealing with this would have been to estimate demand separately for each country, but this is challenging given the relatively short length of the panel.
[26] Appendix B.2 shows that the demand estimates are similar with other natural measures of R_{jt}.
[27] The results are very similar if I use the country's total electricity consumption (from all energy sources) instead.

they increase their demand for LNG for heating, which raises the spot price country j must pay for LNG. This instrument is exogenous as long as changes in temperature in rival countries do not directly affect the LNG demand of country j. The second instrument is electricity consumption from fossil fuels in rival importers. If electricity demand increases in rival importers, their demand for LNG rises, causing the LNG price paid by country j to increase. The instrument is plausibly exogenous provided that country j's idiosyncratic LNG demand shocks, after controlling for its own electricity demand, are uncorrelated with electricity demand in other countries.

The demand estimates are shown in Appendix Table B6. The coefficient on the spot price is negative and statistically significant, and implies a demand elasticity of around -0.92 for the median observation. The coefficients on the demand shifters have the expected sign: LNG demand increases if electricity demand goes up, if the minimum temperature goes down or if the oil price increases. The effective F-statistic for the first stage (following Olea and Pflueger, 2013) is 35, so weak identification is unlikely to be a concern. Appendix B.2 presents additional checks, showing that demand estimates are similar if we explicitly account for import capacity constraints faced by buyers, and are robust to other modelling choices.

5.2 Estimation of Cost Parameters

The starting point for the estimation of the cost parameters is the first-order condition (2). Since shipping costs are known, the only object to estimate is the production cost, or $C(q_{it}, K_{it})$. I assume firms face "soft" capacity constraints (Besanko and Doraszelski, 2004):

$$C(q_{it}, K_{it}) = \delta \frac{1}{1+\nu} \left(\frac{q_{it}}{K_{it}} \right)^{\nu} q_{it} \tag{11}$$

The marginal cost of producing q_{it} units is equal to $\delta \left(\frac{q_{it}}{K_{it}} \right)^{\nu}$. Marginal cost therefore increases with capacity utilization $\frac{q_{it}}{K_{it}}$, with the parameters δ and ν governing the rate at which marginal costs rise with capacity utilization. If $\nu > 1$, marginal cost is strictly convex in capacity utilization.

I estimate the cost function parameters via non-linear least squares. For each parameter guess, I solve numerically for the sellers' predicted spot sales to each buyer, S_{ijt}, via a fixed point routine described in Appendix B.1. The estimator then minimizes the sum of squared deviations between the model predicted spot sales and observed spot sales.

Appendix Table B9 presents cost function estimates. I found it is difficult to identify ν and δ separately from one another.[28] As such I calibrate $\nu = 2$, meaning marginal cost is quadratic in capacity utilization. The estimated δ of 17.3 implies that the marginal cost of production is only \$4.3/MMBtu when a firm operates at 50% capacity utilization, well below the average price of spot LNG in most time periods. By contrast, for a firm operating at 100% capacity utilization, the

[28]This is because high levels of capacity utilization can be rationalized both by high ν and high δ.

marginal cost of production is \$17.3/MMBtu, which is higher than the average spot price. This indicates that it is relatively inexpensive to raise production when capacity utilization is low, but very costly to do so when capacity utilization is high.[29]

5.3 Estimation of contracting and investment model

I now discuss how to estimate the contracting and investment model. The contract premium terms in Stages 3 and 1, and the investment cost, are parametrized as follows:

$$\omega_j^3(q_{ij}, \eta_{ij}^3) = \theta_3 q_{ij} + \frac{\kappa_3}{2} q_{ij}^2 + \eta_{ij}^3 q_{ij} \tag{12}$$

$$\omega_j^1(q_{ij}, \eta_{ij}^1) = \theta_1 q_{ij} + \frac{\kappa_1}{2} q_{ij}^2 + \eta_{ij}^1 q_{ij} \tag{13}$$

$$\Gamma_i(K_i, \eta_i^2) = \gamma_1 K_i + \frac{\gamma_2}{2} K_i^2 + \eta_i^2 K_i \tag{14}$$

I assume the contract premium in both stages is quadratic in the contract quantity q. θ_3 is the buyer's marginal willingness-to-pay for an additional unit of contracted LNG at Stage 3 (relative to spot purchases) when the contract quantity q is zero. κ_3 determines how this marginal willingness-to-pay differs as q increases. θ_1 and κ_1 are defined analogously for Stage 1 contracts. I allow the contract premium to differ between Stage 1 and 3 because buyer preferences for contracts could depend on when the contract is signed. For example, if buyers are motivated by supply assurance motives, and prefer to lock in purchases earlier rather than latter, θ_1 might be higher than θ_3. Alternatively, θ_1 might be lower than θ_3 if buyers prefer not to sign contracts too far in advance.[30]

The investment cost is also a quadratic function of investment K, where γ_1 represents the marginal cost of investment at $K = 0$ while γ_2 determines how this marginal cost varies with the level of investment. Finally, I assume in the baseline specification that the bargaining weight equals τ for all seller-buyer pairs. In robustness checks, I test for differences in τ across buyers and sellers.

The structural parameters to be estimated are therefore θ_3 and κ_3 (contract premium in Stage 3), θ_1 and κ_1 (contract premium in Stage 1), γ_1 and γ_2 (investment cost parameters) and τ (bargaining weights). The estimation of these parameters utilizes the first-order conditions of each of the three stages of the game, namely equation (5) (Stage 3 FOC), equation (8) (Stage 2 FOC), and equation (9) (Stage 1 FOC). I estimate the parameters using non-linear least squares.

Approximation of buyer and seller payoffs: In order to estimate the model, for each parameter guess, we need to be able to solve for the players' payoffs (and derivatives of these payoffs) at each stage of the game. Because of the complex nature of the spot market equilibrium, however, analytical expressions for per-period payoff functions, $\pi_{it}^s(q_t^c, K_t)$ and $\pi_{jt}^b(q_t^c, K_t)$ do not exist. This means

[29] In robustness analyses, reported in Appendix B.3, I check how the estimated costs depend on the assumed nature of capacity constraints, finding similar results and model fit.

[30] The average Stage 1 contract is signed 5.3 years before the start date of the contract (when deliveries begin), compared to 2.8 years for the average Stage 3 contract.

that constructing the exact payoffs is computationally very demanding, since it requires numerically solving the spot market equilibrium for every year while integrating out demand shocks.

Instead, I use parametric approximations of the per-period expected payoff functions when estimating the model, similar to Sweeting (2013) and Barwick and Pathak (2015). I assume that each seller's per-period payoff can be approximated by a set of L_s basis functions $u_1,...,u_{L_s}$, and each buyer's per-period payoff by a set of L_b basis functions $\phi_1,...,\phi_{L_b}$:

$$\pi^s_{it}(q^c_t,K_t) \simeq \sum_{l=1}^{L_s} b^s_l u_l(q^c_t,K_t,x_t);\ \pi^b_{jt}(q^c_t,K_t) \simeq \sum_{l=1}^{L_b} b^b_l \phi_l(q^c_t,K_t,x_t)$$

where b^s_l and b^b_l are unknown approximating parameters that need to be estimated. x_t is a vector of exogenous seller and buyer characteristics, as well as the distance between each seller and buyer.

The state space is a high-dimensional object, since it includes q^c_t (a vector of every contract active in period t) and a K_t (a vector of the capacity of every seller active in period t). To further lessen the computational burden, I assume that instead of keeping track of the state variables of each of their rivals, firms only keep track of two sufficient statistics: the total capacity of rivals, and the total contract quantity signed by all rival firms. This approach has similarities to the notion of oblivious equilibrium developed by Weintraub et al. (2008) and Benkard et al. (2015).

To estimate the approximating parameters, I simulate the spot market model for a large set of random draws of q^c_t, K_t, and ε_t. For each random draw, I solve for the spot market model, integrate over the demand shocks, and derive the per-period expected payoffs to sellers and buyers. In the resulting simulated sample, I then regress the expected payoffs on the basis functions to derive the approximating parameters. In practice, I found that a quadratic approximation works well, delivering a R^2 of 0.98 in the simulated sample; using higher-order polynomials did not significantly reduce the approximation error. Appendix E.1 provides further details on the approximations, including the exact basis functions used.

Identification: The data used for estimation includes contract quantities signed between buyers and sellers in Stage 1 (ex-ante) and Stage 3 (ex-post), as well as capacity investments made by the seller in Stage 2. If contract prices were also observed, they would directly be informative about how the surplus is split between buyers and sellers, and thus pin down the bargaining weight τ. However, contract prices are unobserved.

Instead, τ is identified from variation in the disagreement payoffs of sellers and buyers. Recall from Section 4.4 that as the seller's disagreement payoff worsens, the risk of under-investment increases, and so the seller and buyer will sign larger ex-ante contracts to forestall under-investment. The sensitivity of investment and contracting decisions to the seller's disagreement payoff depends on the Nash bargaining weight τ. If τ is close to 1, then the seller can capture most of the surplus from trade, and changes in the seller's disagreement payoff will have little effect on investment and contracting decisions. If instead τ is close to 0, the seller's disagreement payoff will have a much bigger effect on investment and contracting decisions.[31] Thus, the greater τ is, the *less* sensitive investment and ex-ante contract quantities are to the seller's disagreement payoff. Variation in the buyer's disagreement payoff is likewise helpful for identifying τ: the greater τ is, the *more* sensitive investment and ex-ante contract quantities are to the buyer's disagreement payoff.

In this empirical context, there is excellent variation in the disagreement payoffs for sellers and buyers, both across space and over time. Sellers who are located close to multiple buyers enjoy higher disagreement payoffs than sellers that are far away from most buyers, since they incur lower shipping costs and thus their expected payoff from the spot market is higher. In the same vein, buyer disagreement payoffs also vary as a function of geography. Seller and buyer disagreement payoffs also vary over time, as new sellers and buyers enter the market, as new capacity is built and as existing contracts expire (freeing up more capacity for the spot market). This variation aids with identification despite the absence of systematic data on negotiated prices.

The identification of the investment cost and contract premium parameters is more standard (provided, of course, there is enough variation in disagreement payoffs to identify τ). Intuitively, variation in the seller's value of investing (for example, as the availability of buyers changes over time) helps identify the investment cost parameters. The Stage 1 contracting moments identify θ_1 and κ_1, the contract premium parameters in Stage 1. Similarly, the Stage 3 contracting moments identify the corresponding contract premium parameters θ_3 and κ_3.

Results: Table 3 shows the estimated parameters. Based on the investment cost parameter estimates (γ_1 and γ_2), the average cost of building a 5 mtpa project (which is a median-sized project) is \$4.36 billion for every mtpa of capacity built. γ_2 is estimated to be around -0.28, suggesting there are (modest) economies of scale in building LNG projects: the largest plant (with size 16.5 mtpa) is estimated to have an average cost of \$4.30 billion/mtpa, compared to \$4.41 billion/mtpa for the smallest plant (with size 0.5 mtpa).

The estimated average investment cost (of \$4.36 billion/mtpa) is larger than accounting estimates of the average cost of building liquefaction capacity, which is \$2.7 billion/mtpa.[32] This may be because the accounting cost does not include the cost of financing (e.g., the costs associated with obtaining debt finance), which may be substantial for LNG projects. The accounting cost also ignores other capital costs that sellers may incur (in addition to the cost of building the liquefaction plant), such as the cost of developing upstream infrastructure needed to ensure natural gas is delivered to the plant. Finally, the production cost function does not include any fixed costs of operating and maintaining plants, which instead might be subsumed in the cost of investment that I estimate.

[31] This intuition is clearest in the polar cases where τ is either 1 or 0. For instance, when τ is 1, the buyer only need be paid their disagreement payoff, and the seller captures all the remaining surplus. In this case, changes in the seller's disagreement payoff have no effect on investment (and therefore no effect on ex-ante contracting).

[32] The latter estimate is based on 13 plants for which published estimates of the final cost of investment exist.

Table 3: Contracting and investment parameter estimates

	Estimate	S.E.		Estimate	S.E.
Investment cost, γ_1	89.60	(3.52)	Contract premium (ex-ante) : κ_1	-0.008	(0.011)
Investment cost, γ_2	-0.28	(0.35)	Contract premium (ex-ante) : θ_1	0.22	(0.46)
Bargaining weight, τ	0.64	(0.09)	Contract premium (ex-post): κ_3	-0.009	(0.006)
			Contract premium (ex-post): θ_3	1.40	(0.12)
Number of contracts (Stage 1)	123		Number of contracts (Stage 3)	172	
Number of investments	54				

Note: All parameters estimated using non-linear least squares. Standard errors are heteroskedasticity-robust.

The mean contract premium parameter in Stage 3 (θ_3) equals 1.40 (and is statistically different from 0), while κ_3 is negative implying that the contract premium diminishes as the contract quantity increases. Together, these parameters imply that buyers are on average willing to pay a per-unit premium of \$1.34/MMBtu for LNG purchased under *ex-post* long-term contracts as opposed to spot purchases, which is about 15% of the average spot price (which is \$8.7/MMBtu). By contrast, the evidence for a contract premium in Stage 1 is weaker. The mean of the Stage 1 contract premium (θ_1) is 0.22, and statistically indistinguishable from 0. Taking into account diminishing returns to contracting (since κ_1 is negative), the average per-unit contract premium for *ex-ante* contracted purchases is only \$0.12/MMbtu, or 1.4% of the average spot price. The finding that a contract premium exists for ex-post but not ex-ante contracting is consistent with buyers having a preference for signing long-term contracts (due to, for example, supply assurance or transaction cost avoidance motives), but not wishing to commit to long-term contracts too far in advance. Appendix B.4 presents additional estimates investigating the determinants of the contract premium, finding that the contract premium is increasing in the ability of countries to enforce contracts (see Table B12).

The bargaining weight τ is estimated to equal 0.64. Thus, although sellers are able to capture a greater share of the surplus from trade than buyers, buyers still have considerable bargaining power. We can reject the hypothesis that sellers have the ability to make take-it-or-leave-it offers to the buyer. This echoes much of the empirical Nash bargaining literature that generally finds firm behavior is inconsistent with the take-it-or-leave-it bargaining model (Crawford and Yurukoglu, 2012; Grennan, 2013; Ho and Lee, 2017). In Appendix B.4, Table B11, I explore specifications allowing heterogeneity in bargaining power, finding little evidence that bargaining power differs significantly across sellers and buyers. As such, I use the baseline specification with a single bargaining weight τ for the remainder of the analysis.

Sellers' incentives to invest are dampened by buyer bargaining power. If buyers instead had no bargaining leverage (i.e., if the bargaining weight were 1 and if the buyers had no outside option of going to the spot market), the seller's marginal benefit of investing would go up by

22% on average. This creates an incentive for sellers and buyers to sign larger ex-ante contracts to forestall under-investment, explaining why ex-ante contracting persists despite buyers having a higher willingness-to-pay for ex-post contracts (as shown by the contract premium being higher in Stage 3 than in Stage 1). In the counter-factual analysis that follows, I further investigate the role of ex-ante contracting in mitigating under-investment by estimating how investment would adjust if sellers and buyers were not able to sign ex-ante contracts.

While long-term contracts are beneficial to the contracting parties, they may also exert externalities on rival sellers and buyers. For each contract I compute the *marginal externality*, or the derivative of the total welfare of non-contracting parties with respect to the contract quantity. If contracts exerted no externality, this quantity would equal 0 on average. Instead, I find that the marginal externality is on average equal to -\$0.91/MMBtu, implying that the marginal external cost of contracting is equal to about 10.4% of the average price of LNG on the spot market. Further decomposing this between buyers and sellers, I find that contracts exert a small positive externality on rival sellers (on average \$0.26/MMBtu), but a sizeable negative externality on rival buyers (on average -\$1.17/MMBtu), intuitively because a contract signed by one buyer reduces the average supply of LNG available to other buyers. The existence of negative externalities suggests there is likely to be over-contracting in equilibrium.

Finally, as a validation of the empirical methodology of estimating bargaining power without observing negotiated prices, I compare the contract prices predicted by the model with contract prices computed from customs data (which are available for China, Japan and Korea). Most LNG contracts use relatively simple formulas linking the LNG price to an oil price index; thus, customs data can be used to estimate the statistical relationship between LNG prices and oil prices, which provides information on the contract price formulas (Agerton, 2017). Appendix Figure E5 shows that the model-predicted contract prices are quite similar to the contract prices inferred from customs data (with a correlation of 0.45), despite the fact that the latter data is not utilized in estimation.

6 Counter-factual Analysis

Using the estimated model, I carry out several counter-factual exercises in order to explore the consequences of using long-term contracts for investment, allocations and welfare.

Solving the full multi-stage game is computationally intensive, especially for projects where a seller contracts with multiple buyers in both Stages 1 and 3. To reduce the computational burden, I assume in all the simulations that for any given investment project, a seller can only contract with a single buyer.[33] While this is a strong assumption, I find that the baseline simulations under this assumption yields a similar level of investment to what I see in the data.[34] A drawback of this

[33] In investment projects where a seller contracts with multiple buyers, I select the buyer with the largest contract.

[34] The capacity built in an average export project is 6.59 mtpa in the baseline simulations, versus 6.66 mtpa in the data.

assumption is that it limits the richness of strategic interactions that are permitted in the counter-factual simulations, since I do not allow a seller to bargain simultaneously with multiple buyers.

6.1 Contracts and investment

I begin by quantifying how the ability to sign long-term contracts affects investment. I solve for sellers' investment decisions under two counter-factual contracting regimes: (i) "no contracting", where sellers and buyers cannot sign any long-term contracts, and can only trade on the spot market (ii) "no ex-ante contracting", where sellers and buyers cannot sign ex-ante contracts in Stage 1, but can sign contracts ex-post in Stage 3 or trade on the spot market. I compare these to a baseline regime where sellers and buyers face no restrictions on either ex-ante or ex-post contracting.

The counter-factuals reported in this subsection are "partial equilibrium" counter-factuals, where I solve for the investment and contracting choices of the agents involved in each investment project, while holding fixed the investment and contracting choices of the rest of the industry participants. I purposely abstract from general equilibrium considerations for now, since my goal here is to investigate how investment incentives at the plant level are affected by bargaining power and the ability to sign long-term contracts. I explore the general equilibrium effects of contracting in Section 6.3.

I find that that if sellers are not permitted to sign long-term contracts with buyers, they would lower investment by 27% on average (top panel of Table 4). The reduction in investment is primarily due to the inability to sign ex-ante contracts: if sellers and buyers can sign ex-post (but not ex-ante) contracts, investment would still decline by 24%.

Table 4: Average capacity of project (mtpa), with restrictions on long-term contracting

	Benchmark	No contracting, PE		No ex-ante contracting, PE	
			% change		% change
All projects	6.59	4.78	*-27.4%*	5.01	*-24.0%*
Distance from nearby buyers					
Greater than median	7.48	5.01	*-33.0%*	5.27	*-29.5%*
Lower than median	5.13	4.34	*-15.3%*	4.56	*-11.0%*
Plant Capacity					
Greater than median	9.90	7.03	*-29.0%*	7.36	*-25.7%*
Lower than median	3.28	2.53	*-22.8%*	2.66	*-19.0%*

Note: In the benchmark regime, sellers and buyers can sign long-term contracts in either Stage 1 or Stage 3. In the "no-contracting" regime, no contracting is permitted. In the "no ex-ante contracting" regime, sellers and buyers can only sign contracts in Stage 3. The table shows the average capacity built by sellers, and (in italics) the percentage change relative to the benchmark regime. All counter-factuals are partial equilibrium (PE).

The remaining panels of Table 4 show how investment in each of these counter-factual regimes varies by seller characteristics. Sellers who are geographically more isolated than the median seller, and who thus have less favorable outside options from trading on the spot market, reduce their investment by 29.5% when they cannot sign ex-ante long-term contracts, compared to 11% for sellers located closer to buyers than the median seller.[35] In a similar vein, the reduction in investment from not being able to contract ex-ante is larger for sellers building large plants (26%) than for sellers building small plants (19%). Sellers making bigger investments have a weaker outside option when negotiating ex-post contracts with buyers, since they are left with a larger excess capacity to offload on the spot market in the event of disagreement.[36] This makes them more reliant on ex-ante long-term contracting to mitigate under-investment.

6.2 Contracts and allocative efficiency

Next, I quantify the allocative efficiency consequences of using long-term contracts. To do so, I simulate the industry equilibrium with and without long-term contracts. In order to cleanly quantify the allocative efficiency effects, I abstract from any investment considerations for now, by holding the investment of every seller fixed in each counter-factual. In the next section (Section 6.3), I consider investment and allocation effects jointly.

Even ignoring investment considerations, theory is ambiguous on whether long-term contracts improve or worsen allocative efficiency. As discussed in Section 4.4 (page 25), long-term contracts reduce the flexibility of sellers in meeting demand shocks, but decrease distortions from market power in the spot market. I find that the former effect dominates, so that switching from long-term contracting to spot trade results in sizeable allocative efficiency gains. Aggregate welfare (discounted back to 2001) would increase by \$38 bn in the no-contracting regime relative to the benchmark regime, or around 0.7% of overall welfare (top panel of Table 5). Buyers benefit most, with buyer surplus increasing by 1.6% if long-term contracting were eliminated.

These allocative efficiency gains largely arise because trading on the spot market allows sellers to respond more efficiently to demand shocks: as the second panel of Table 5 shows, if we were to eliminate demand uncertainty, the allocative efficiency gains would become far smaller, decreasing from 0.7% to only 0.1%.

The flexibility gains from reducing contract usage are most significant when a large share of seller capacity is tied up under long-term contracts (as is the case in practice), leaving very little spare capacity on the spot market that can be deployed to deal with demand fluctuations. As the share of capacity that is contracted decreases, though, the flexibility gain from reducing contract usage further diminishes in size, so that at some point the pro-competitive effect of long-term

[35] As discussed in Section 3, I use the 25th percentile of the distance from the seller to buyers as a measure of the geographic isolation of the seller. The results are very similar if we instead use the mean distance from the seller to all buyers.

[36] Williamson (1983) describes this as the “dedicated assets" problem, where a seller makes a large investment for the primary purpose of selling a large quantity to a single buyer, and would end up with significant excess capacity if that buyer did not end up purchasing from them.

Table 5: Allocative efficiency with and without long-term contracting, holding investment fixed

	Benchmark	No contracting	% change
Baseline assumptions			
Welfare ($ bn)	5,469	5,507	*0.7%*
Seller surplus ($ bn)	2,535	2,527	*-0.3%*
Buyer surplus ($ bn)	2,934	2,980	*1.6%*
No demand uncertainty			
Welfare ($ bn)	5,298	5,306	*0.1%*
Seller surplus ($ bn)	2,556	2,543	*-0.5%*
Buyer surplus ($ bn)	2,743	2,762	*0.7%*

Note: In the benchmark regime, sellers and buyers can sign long-term contracts in either Stage 1 or Stage 3. In the "no-contracting" regime, no contracting is permitted. Investment is held fixed (at observed levels) in all simulations. The table shows discounted total welfare, total seller surplus and total buyer surplus in US$bn from 2001 onwards.

contracts begins to dominate. Appendix C.2 shows that this happens once the aggregate volume of contracting is reduced to approximately 25% of the baseline volume; beyond that point, further reductions in long-term contract usage lead to *lower* allocative efficiency, due to the Allaz and Vila (1993) market power effect. Thus, allocative efficiency is maximized by a mix of spot trading and long-term contracting (rather than trading solely on the spot market), albeit one that is much more heavily skewed towards spot trading than what we observe in practice.

To gain further insight into the allocative efficiency effects of using long-term contracts, I consider the impact of a hypothetical shutdown of Russian natural gas exports to Europe, leading to a large increase in European demand for LNG. I investigate how firms in the LNG industry respond to this demand shock with and without long-term contracts (continuing to hold capacity fixed).[37]

I find the industry responds more efficiently to the demand shock if long-term contracts were not used. In the event of a hypothetical shutdown of Russian natural gas exports to Europe, I estimate that European LNG imports in the baseline regime increase only by 46.8 mt annually, versus 49.8 mt in the no-contracting regime (Appendix Table C15). This leads to an annual loss of industry surplus of around $1.7 bn. Long-term contracts result in an inefficiently muted response to the demand shock, since some sellers bound by long-term contracts do not re-allocate LNG to the buyers in Europe experiencing the demand shock.

This demand-shock counter-factual also illustrates why buyers have unilateral incentives to sign long-term contracts, despite their lack of flexibility; and why contracts exert negative externalities. In the event of a shutdown of Russian natural gas exports to Europe, European buyers benefit from their own existing long-term contracts, but are negatively affected by contracts signed by other buyers (Appendix C.3). More generally, a buyer's own long-term contracts are especially valuable

[37] In Appendix C.3, I also study the LNG industry response to the Fukushima nuclear disaster in Japan.

during a demand shock, since they protect them against sellers' exercise of market power (this is the insight of Allaz and Vila, 1993); but these same contracts hurt other buyers when *they* experience demand shocks, by reducing the spare supply of LNG that can be used to met their demand shocks. These forces create strong unilateral incentives to sign long-term contracts, even if these contracts (on aggregate) exert negative externalities on other parties. Negative contracting externalities may lead to over-contracting in equilibrium, which I explore next.

6.3 Welfare effects of using long-term contracts

The counter-factuals so far have examined the investment and allocative efficiency effects of using long-term contracts in isolation. To study the welfare effects of long-term contracting, though, we need to consider how these forces interact with each other in equilibrium. I therefore simulate a counter-factual regime where sellers and buyers cannot sign any long-term contracts after a specified date (which I choose to be 1995). Unlike the partial equilibrium counter-factuals considered earlier, I now account for general equilibrium considerations. This means that each seller, when investing, takes into account how other sellers are going to adjust their investment levels in response to not being allowed to contract. Solving the full industry equilibrium is computationally involved and the Gauss-Jacobi fixed point algorithm for doing so is described further in Appendix C.1.

Table 6: Welfare effects of long-term contracting, in general equilibrium

	Benchmark	No contracting	% change
Welfare ($ bn)	5,469	5,493	*0.4%*
Seller surplus ($ bn)	2,535	2,579	*1.7%*
Buyer surplus ($ bn)	2,934	2,914	*-0.7%*
Average capacity (mtpa)	6.59	5.89	*-10.6%*

Note: In the benchmark regime, sellers and buyers can sign long-term contracts in either Stage 1 or Stage 3. In the no-contracting regime, I solve for the new industry equilibrium when no contracting is permitted, allowing sellers to adjust their investment levels. The first three rows show discounted total welfare, total seller surplus and total buyer surplus in US$bn from 2001 onwards. The fourth row shows the average capacity of each export project built.

The results are shown in Table 6. We saw earlier that when the seller and buyers in an individual project cannot sign long-term contracts (but the rest of the industry can), sellers reduce investment by 27.4%. When *no* sellers and buyers can sign long-term contracts, though, sellers only reduce investment by 10.6% (Table 6). Thus, the under-investment from not being able to sign long-term contracts is much less severe in general equilibrium. This is because when no one in the industry can sign long-term contracts, each seller realizes that other sellers, facing hold-up risk, will under-invest. The reduced investment by other sellers means the seller can expect to receive a higher price on the spot market; this partially counteracts the negative effect of the seller not being able to sign their own long-term contracts.

Eliminating long-term contracting thus results in sizeable welfare gains equal to $23 bn, or about 0.4% of total welfare: the under-investment by sellers is more than compensated for by gains in allocative efficiency. These welfare gains are unevenly distributed, with sellers' surplus increasing by 1.7%, while buyers' surplus decrease by 0.7%. This is partly because eliminating contracting also has the effect of reducing competition on the spot market, in a reversal of the Allaz and Vila (1993) effect. Furthermore, the reduction in aggregate investment also serves to reduce competition and raise spot prices, which benefits sellers at the expense of buyers.

There are welfare gains from restricting long-term contracts (in spite of their voluntary nature) because of contracting externalities. Although there are allocative efficiency benefits from reducing the use of long-term contracts, many of these benefits accrue to third parties (e.g., other buyers of LNG). The buyers and sellers who are actually involved in the contract negotiation do not fully internalize these benefits, and have strong private incentives to sign long-term contracts, both as a way to induce the seller to invest more, and to protect the buyer against ex-post demand shocks. This results in over-contracting in equilibrium. The presence of these externalities helps explain why an inefficiently high degree of long-term contracting can persist and why the growth of the LNG spot market has been relatively slow.

Thus, policies that restrict the use of long-term contracts are potentially welfare-enhancing. However, such policies are challenging to implement, given the lack of a regulatory body with the power to implement industry-wide policies across the global LNG industry. Limits on long-term contracts would be effective only with the cooperation of all sellers and buyers, yet individual sellers and buyers would have strong incentives to renege and sign long-term agreements. This may explain the absence of any regulatory attempts to limit the use of long-term LNG contracts.

Finally, it is important to note that eliminating long-term contracts, even if it increases welfare, is unlikely to be the first-best policy that maximizes welfare. As we saw in Section 6.2, even ignoring investment considerations, it would still be optimal to allow some long-term contracting (roughly 25% of the current volume) due to the pro-competitive benefits of contracts. Once investment considerations are taken into account, the socially optimal level of contracting is higher still, since long-term contracts have considerable value in reducing under-investment. Thus, the first-best organizational structure for the industry is very likely to involve some mix of long-term contracts and spot trade (albeit one that features considerably less contracting than what we see in practice).

6.4 Effects of banning resale restrictions

The inflexibility of long-term contracts is partly by design: contract quantities are difficult to adjust ex-post due to the widespread prevalence of various types of resale restrictions (such as destination and diversion clauses). So instead of restricting whether or not parties can sign long-term contracts, the preferred approach of regulators in recent years has been to limit the use of resale restrictions in long-term contracts (while still permitting parties to freely sign contracts that do not include such restrictions). Both in Europe and in Japan, anti-trust regulators have attempted to prohibit the use of destination clauses in LNG contracts. Such policies, however, have not yet been universally implemented, and the impact of these policies is still not fully understood. In this section, I carry out a counter-factual that assesses the long-run consequences of the removal of all resale restrictions from LNG contracts.

In a world without resale restrictions, buyers of contracted LNG face substantially lower costs of re-selling LNG; this increases flexibility since buyers can re-sell LNG in response to unanticipated demand shocks. Furthermore, contracted buyers can also arbitrage away inter-regional price differentials, substantially reducing sellers' ability to engage in spatial price discrimination; this reduces the distortion from market power. However, by reducing sellers' market power, the policy lowers the share of profits accruing to sellers and increases the bargaining leverage of buyers, worsening under-investment. The dilemma faced by regulators is therefore that the removal of resale restrictions improves allocative efficiency but may potentially reduce investment incentives.[38]

I model the removal of resale restrictions as follows. Because contracted buyers can now engage in resales, sellers now face the threat of arbitrage on the spot market. If a seller were to price discriminate and charge different spot prices (net of shipping costs) to different buyers, then buyers of contracted LNG can arbitrage away the price differences. I assume this threat of arbitrage is so strong that sellers are forced to behave competitively and are unable to exercise any market power.

The assumption of competitive behavior rules out the possibility that sellers might continue to exercise market power even if they are unable to price discriminate. It also assumes away the possibility that the arbitrageurs (buyers) might themselves exercise market power and price discriminate when reselling LNG; however this possibility is rather remote since there are many more contracted buyers than sellers and the market would be much less concentrated when buyers can also act as re-sellers. Finally, I implicitly assume there cannot be any costs of arbitrage (e.g., transaction costs). If one or more of the above conditions are violated, market power might continue to be significant even when resale restrictions are removed; in that case, the actual welfare gains from the policy may be lower than what I find.

Under this assumption, the allocation of LNG will be competitive, regardless of long-term contracts. Even if a seller i and j agrees to a contract with quantity q^c_{ij}, the actual sale of LNG from i to j will be the competitive allocation q^*_{ij}, where q^*_{ij} may differ from q^c_{ij}. The contracts do, however, affect the payoffs to the buyers and sellers. If a seller delivers less LNG to the buyer than promised (that is, if $q^*_{ij} < q^c_{ij}$), then the seller refunds buyer j for the shortfall in deliveries, with the refund

[38] A similar dilemma arises in the regulation of the trade of pharmaceutical drugs, where permitting parallel trade has the benefit of preventing drug manufacturers from price discriminating across buyers in different countries, but may lower the profits of manufacturers and reduce their incentives to innovate (Dubois and Sæthre, 2020).

price exactly equal to the prevailing spot price faced by buyer j.[39] If instead the seller delivers more LNG to the buyer than promised (i.e., $q^*_{ij} > q^c_{ij}$), then the buyer only has to pay the spot price for the additional LNG delivered beyond the contracted amount (since the contracted amount is already paid for).[40] Thus long-term contracts can still be used to reduce potential under-investment, or for hedging purposes, but no longer restrict the flexibility of sellers in allocating LNG.

I assume that the policy takes effect in the year 2012. I find that the policy leads to a sizeable decrease in investment. As Table 7 shows, the removal of resale restrictions reduces sellers' market power on the spot market, so the average spot price decreases by 17%. This directly reduces sellers payoffs from spot sales. It also worsens the bargaining leverage of sellers when negotiating long-term contracts with buyers, since the seller has a weaker outside option and the buyer has a stronger outside option (given that spot prices are lower). As a result, the average contract price decreases by 6.8%. Sellers and buyers, anticipating the greater risk of under-investment, partly compensate for this by signing larger contracts (to insulate the seller to some extent from the lower spot price), with the average size of a contract increasing by almost 70%: it is less costly to sign large contracts in this environment, since any excess quantity contracted can always be re-sold. In spite of this, sellers still have weaker incentives to invest, so that investment decreases by 15.8%.

Table 7: Investment and welfare impact of prohibiting resale restrictions

	Benchmark	No resale restrictions	% change
Welfare ($ bn)	5,469	5,982	*9.4%*
Seller surplus ($ bn)	2,535	1,942	*-23.4%*
Buyer surplus ($ bn)	2,934	4,040	*37.7%*
Average capacity of plant (mtpa)	6.26	5.27	*-15.8%*
Average size of contract (mtpa)	2.39	4.07	*70.2%*
Average contract price ($/MMBtu)	10.55	9.83	*-6.8%*
Average spot price ($/MMBtu)	9.76	8.10	*-17.0%*

Note: In the benchmark regime, sellers and buyers sign long-term contracts that have destination clauses. In the "no resale restrictions" regime, the use of destination clauses (and other resale restrictions) is banned in 2012, so that from then on all buyers under long-term contracts have the option of re-selling LNG. The first three rows of the table shows discounted total welfare, total seller surplus and total buyer surplus in US$bn from 2001 onwards. The fourth row shows the average capacity of each export project built after 2012, in million tonnes per annum (mtpa). The fifth row reports the average size of contracts signed by sellers and buyers in these projects, while the fifth and sixth rows show the average contract price and the average spot price (in $/MMBtu).

However, the absence of resale restrictions leads to a substantially more efficient allocation of LNG, both from reduced market power and the absence of contractual rigidities. As such, despite

[39]This ensures that the buyer is fully compensated for the shortfall, since the refund paid to the buyer equals exactly equals the amount the buyer would have to pay to purchase that same amount of LNG on the spot market.

[40]This is similar to how forward contracts in electricity markets are modelled (e.g., see Hortacsu and Puller, 2008).

the reduction in investment, the welfare gains from removing resale restrictions are sizeable, with welfare increasing by around $513 bn (or 9.4%).

The welfare gains from removing resale restrictions are not enjoyed by both buyers and sellers, though: while buyer surplus increases by 37.7%, seller surplus declines by 23.4%. The uneven division of gains between buyers and sellers helps explain why resale restrictions are still widely used in LNG contracting, despite their inefficiency. Unsurprisingly, LNG exporters have mostly been opposed to the removal of destination clauses and other forms of resale restrictions, while LNG importers (such as the EU or Japan) have been at the forefront of recent regulatory attempts to prohibit such restrictions. In the absence of a single industry-wide regulator, these conflicting interests of sellers and buyers make it difficult to prohibit the use of destination clauses, despite the substantial welfare gains from doing so.[41]

7 Conclusion

In markets where firms make large sunk cost investments, ex-post bargaining can reduce the surplus sellers enjoy from the investment, resulting in under-investment. Ex-ante long-term contracts are valuable in these markets for mitigating under-investment. However, rigidities in contract design can inhibit the ability of firms to respond to demand shocks. In this paper, I develop and estimate a structural model of investment, contracting and spot trade to quantify the trade-off firms in the LNG industry face between under-investment and contract rigidity.

The empirical analysis highlights the inefficiencies that can arise from the use of long-term contracts. Specifically, rigid long-term contracts impose negative externalities on other firms, leading to excessive use of long-term contracts in equilibrium. The degree of over-contracting is severe enough that eliminating long-term contracting would increase welfare, in spite of the reduction in investment. Policies that reduce contractual rigidities by prohibiting the use of resale restrictions, which have been promoted by anti-trust authorities in EU and Japan, also reduce investment, but lead to substantial welfare gains amounting to over $500 bn (or more than 9% of total welfare). More broadly, the results of this paper suggest that there may be efficiency gains from regulating the use of long-term contracts in markets where contracting rigidities are important.

References

AGERTON, M. (2017): "Global LNG pricing terms and revisions: An empirical analysis," *The Energy Journal*, 38.

[41] For example, in the LNG contract negotiations between Germany and Qatar in the summer of 2022, a major source of contention was over whether the contract should include destination clauses. See https://www.reuters.com/business/energy/exclusive-germany-qatar-odds-over-terms-talks-lng-supply-deal-sources-2022-05-09/.

ALLAZ, B. AND J.-L. VILA (1993): "Cournot Competition, Forward Markets and Efficiency," *Journal of Economic Theory*, 59, 1–16.

ASON, A. (2019): "Price reviews and arbitrations in Asian LNG markets," Oxford Institute for Energy Studies: NG 144.

BARWICK, P. J. AND P. PATHAK (2015): "The Cost of Free Entry: An Empirical Analysis of Real Estate Agents in Greater Boston," *RAND Journal of Economics*, 46, 103–145.

BENKARD, C. L., P. JEZIORSKI, AND G. Y. WEINTRAUB (2015): "Oblivious equilibrium for concentrated industries," *The RAND Journal of Economics*, 46, 671–708.

BESANKO, D. AND U. DORASZELSKI (2004): "Capacity dynamics and endogenous asymmetries in firm size," *RAND Journal of Economics*, 23–49.

BHATTACHARYA, V. (2021): "An empirical model of R&D procurement contests: An analysis of the DOD SBIR program," *Econometrica*, 89, 2189–2224.

BLOUIN, A. AND R. MACCHIAVELLO (2019): "Strategic default in the international coffee market," *The Quarterly Journal of Economics*, 134, 895–951.

BOLTON, P. AND M. D. WHINSTON (1993): "Incomplete contracts, vertical integration, and supply assurance," *The Review of Economic Studies*, 60, 121–148.

BUSHNELL, J. B., E. T. MANSUR, AND C. SARAVIA (2008): "Vertical Arrangements, Market Structure, and Competition: An Analysis of Restructured US Electricity Markets," *American Economic Review*, 98, 237–266.

CHIPTY, T. AND C. M. SNYDER (1999): "The role of firm size in bilateral bargaining: A study of the cable television industry," *Review of Economics and Statistics*, 81, 326–340.

CRAWFORD, G. S. AND A. YURUKOGLU (2012): "The welfare effects of bundling in multichannel television markets," *American Economic Review*, 102, 643–85.

CROCKER, K. J. AND S. E. MASTEN (1988): "Mitigating contractual hazards: Unilateral options and contract length," *RAND Journal of Economics*, 327–343.

DUBOIS, P. AND M. SÆTHRE (2020): "On the effect of parallel trade on manufacturers' and retailers' profits in the pharmaceutical sector," *Econometrica*, 88, 2503–2545.

EDLIN, A. AND S. REICHELSTEIN (1996): "Holdups, Standard Breach Remedies, and Optimal Investment," *American Economic Review*, 86, 478–501.

EUROPEAN COMMISSION (2017): "Follow-up study to the LNG and storage strategy," .

GOWRISANKARAN, G., A. NEVO, AND R. TOWN (2015): "Mergers when prices are negotiated: Evidence from the hospital industry," *American Economic Review*, 105, 172–203.

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

GROSSMAN, S. J. AND O. D. HART (1986): "The costs and benefits of ownership: A theory of vertical and lateral integration," *Journal of Political Economy*, 94, 691–719.

HARDING, R. AND D. SHEPPARD (2017): "Japan outlaws restrictions on resale of LNG cargoes," *Financial Times*.

HARRIS, A. AND T. M. A. NGUYEN (2022): "Long-term Relationships and the Spot Market: Evidence from US Trucking," Working Paper.

HART, O. AND J. TIROLE (1990): "Vertical Integration and Market Foreclosure," *Brookings Papers on Economic Activity*, 21, 205–286.

HARTLEY, P. (2015): "The Future of Long-term LNG Contracts," *The Energy Journal*, 36.

HO, K. AND R. S. LEE (2017): "Insurer competition in health care markets," *Econometrica*, 85, 379–417.

——— (2019): "Equilibrium provider networks: Bargaining and exclusion in health care markets," *American Economic Review*, 109, 473–522.

HORN, H. AND A. WOLINSKY (1988): "Bilateral monopolies and incentives for merger," *RAND Journal of Economics*, 408–419.

HORTACSU, A. AND S. L. PULLER (2008): "Understanding strategic bidding in multi-unit auctions: a case study of the Texas electricity spot market," *RAND Journal of Economics*, 39, 86–114.

HUBBARD, T. N. (2001): "Contractual form and market thickness in trucking," *RAND Journal of Economics*, 369–386.

IEA (2013): *Developing a Natural Gas Trading Hub in Asia: Obstacles and Opportunities*, Paris: OECD/IEA.

ITO, K. AND M. REGUANT (2016): "Sequential Markets, Market Power, and Arbitrage," *American Economic Review*, 106, 1921–1957.

JOSKOW, P. (1987): "Contract Duration and Relationship-Specific Investments: Empirical Evidence from Coal Markets," *American Economic Review*, 77, 168–185.

KLEIN, B., R. G. CRAWFORD, AND A. A. ALCHIAN (1978): "Vertical integration, appropriable rents, and the competitive contracting process," *The Journal of Law and Economics*, 21, 297–326.

MACCHIAVELLO, R. AND J. MIQUEL-FLORENSA (2017): "Vertical integration and relational contracts: Evidence from the Costa Rica coffee chain," CEPR Discussion Paper No. DP11874.

MACCHIAVELLO, R. AND A. MORJARIA (2015): "The value of relationships: evidence from a supply shock to Kenyan rose exports," *American Economic Review*, 105, 2911–45.

MACKAY, A. (2022): "Contract duration and the costs of market transactions," *American Economic Journal: Microeconomics (Forthcoming)*.

MASTEN, S. E. AND K. J. CROCKER (1985): "Efficient adaptation in long-term contracts: Take-or-pay provisions for natural gas," *American Economic Review*, 75, 1083–1093.

NEUMANN, A., S. RÜSTER, AND C. VON HIRSCHHAUSEN (2015): "Long-term contracts in the natural gas industry: Literature survey and data on 426 contracts (1965-2014)," .

NUNN, N. (2007): "Relationship-specificity, incomplete contracts, and the pattern of trade," *The Quarterly Journal of Economics*, 122, 569–600.

OIES (2017): *The Outlook for Floating Storage and Regasification Units (FSRUs)*, Oxford Institute for Energy Studies.

OLEA, J. L. M. AND C. PFLUEGER (2013): "A robust test for weak instruments," *Journal of Business & Economic Statistics*, 31, 358–369.

RUESTER, S. (2009): "Changing contract structures in the international liquefied natural gas market: a first empirical analysis," *Revue d'économie industrielle*, 127, 89–112.

RYAN, N. (2020): "Contract enforcement and productive efficiency: Evidence from the bidding and renegotiation of power contracts in India," *Econometrica*, 88, 383–424.

——— (2021): "Holding Up Green Energy," National Bureau of Economic Research working paper 29154.

SEGAL, I. (1999): "Contracting with externalities," *The Quarterly Journal of Economics*, 114, 337–388.

SWEETING, A. (2013): "Dynamic Product Positioning in Differentiated Product Markets: The Effect of Fees for Musical Performance Rights on the Commercial Radio Industry," *Econometrica*, 81, 1763–1803.

TALUS, K. (2011): "Long-term natural gas contracts and antitrust law in the European Union and the United States," *Journal of World Energy Law and Business*, 4, 260–315.

TOLVANEN, J., O. DARMOUNI, AND S. ESSIG ABERG (2022): "Pulp Friction: The Value of Quantity Contracts in Decentralized Markets," SSRN working paper 3919592.

WEEMS, P. (2016): "LNG Disputes on the Rise," King & Spalding LLP report.

WEINTRAUB, G. Y., C. L. BENKARD, AND B. VAN ROY (2008): "Markov perfect industry dynamics with many firms," *Econometrica*, 76, 1375–1411.

WILLIAMSON, O. E. (1975): "Markets and hierarchies," *New York*, 2630.

——— (1983): "Credible commitments: Using hostages to support exchange," *American Economic Review*, 73, 519–540.

WILSON, J. D. (2012): "Chinese resource security policies and the restructuring of the Asia-Pacific iron ore market," *Resources Policy*, 37, 331–339.

ZAHUR, N. B. (2022): "Market Power and Spatial Price Discrimination in the Liquefied Natural Gas Industry," Working Paper.

Online Appendix for “Long-term contracts and efficiency in the liquefied natural gas industry”

Nahim Bin Zahur

December 2022

A Additional Industry Details and Descriptive Evidence

This section provides additional details about the industry, data and descriptive evidence.

A.1 Industry Details and Summary Statistics

Data on spot prices and shipping costs: Data on weekly LNG spot prices and shipping costs is compiled from several sources. The most comprehensive of these datasets comes from Waterborne Energy.[42] 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.[43] 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.

Industry Details and Summary Statistics: This paper focuses on the LNG market up until 2017, when derivatives trade played only a very limited role in the market. Historically, financial markets in LNG have been very limited in size and scope, in stark contrast to the global crude oil market or the domestic natural gas market in the United States. Since 2017, though, there has been increasing trade in LNG derivatives. Derivatives represented only about 2% of LNG trade volumes at the beginning of 2017, but by the end of 2018, the share had grown to around 23% (Stapczynski and Murtagh, 2019). The size of the derivatives market is still small relative to the physical market: for comparison, in the crude oil market, derivatives volumes account for around 17 times the volume of physical trade (Terazono, 2019).

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

[43] 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.

Because of the large capital costs involved in the construction of a liquefaction terminal, LNG export projects are typically joint ventures between multiple firms, with the median project having 4 project partners. LNG buyers sometimes purchase small equity stakes in export projects: across all export projects built after 1995, the average equity share of buyers was 8.2%. However, full vertical integration, where the same firm controls the entire supply chain, is very rare in the LNG industry. This is partly as many exporting countries require that international joint ventures for LNG be either fully or majority owned by domestic firms, making full vertical integration infeasible.[44]

Table A1 contains summary statistics on key variables used in the analysis. The first two panels include trade flows and shipping costs (defined at the exporter-importer-year level); export and spot prices and total importers (defined at the importer-year level). Panel C show key statistics for the dataset of 464 long-term contracts. The average long-term contract is 17 years in duration, and is signed 3.6 years before the start date of deliveries. Panel D summarizes the data on export projects. Export projects are generally very large in size, with the typical investment equal to 6.94 mtpa (for context, the average export capacity of each LNG exporting country is 14 mtpa). Time-to-build is substantial: on average 4.3 years pass between the time when a FID is announced and the time when the export project begins operating.

Table A1: Summary Statistics

	Obs.	Mean	S.D.	Min	Max
Panel A. Exporter-Importer-Year					
Spot Trade (mt)	6,406	0.10	0.39	0	7.70
Contracted Trade (mt)	6,406	0.35	1.47	0	23.90
Shipping Cost (US$/MMBtu)	9,812	1.30	0.87	0.06	5.08
Panel B. Importer-Year					
Total Imports (mt)	359	8.24	15.45	0	89.19
Spot Prices (US$/MMBtu)	317	8.65	3.65	2.52	16.59
Panel C. Contract-level					
Annual contract quantity (mtpa)	464	1.28	1.10	0.04	5.2
Duration (years)	464	17.23	6.41	4	42
Time from signature to start date (years)	464	3.61	2.17	0	12
Signature Year	464	2004	11.31	1963	2018
Panel D. Export project-level					
Capacity (mtpa)	74	6.94	5.02	0.5	28.9
Year of Final Investment Decision	74	2003	14.39	1959	2021
Time from FID to start date (years)	74	4.30	1.32	2	9.08

Note: All trade variables (spot trade, contracted trade, total exports, etc.) are measured in million tonnes or mt. Capacity and annual contract quantity are measured in million tonnes per annum, or mtpa. Spot prices and shipping costs are measured in US$/MMBtu.

[44] This is the case, for example, in Qatar and Indonesia, historically two of the world's largest LNG exporters.

Figure A1 shows the evolution of liquefaction (export) capacity, regasification (import) capacity and LNG trade over time. It illustrates that capacity utilization for exporters is high, whereas there is substantial amounts of excess import capacity.

Figure A1: Liquefaction capacity, regasification capacity and LNG trade over time

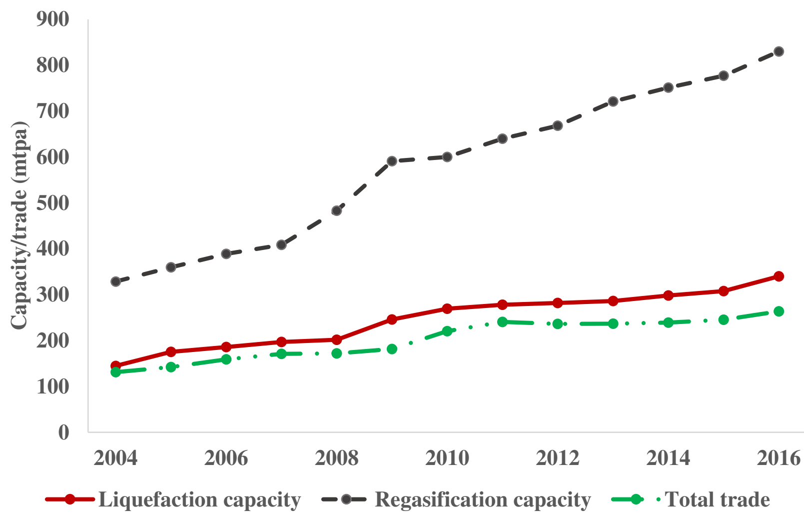

Note: The figure plots global liquefaction capacity, regasification capacity and the volume of LNG trade over time. Liquefaction capacity refers to the annual nameplate capacity of liquefaction projects (which are used to export LNG). Regasification capacity refers to the annual nameplate capacity of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

Figure A2 shows LNG imports to different regions, broken down by long-term contracts (Figure A2a) and short-term and spot trade (Figure A2b).

Figure A2: LNG imported by different regions

(a) LNG imports under long-term contracts **(b)** Short-run and spot LNG imports

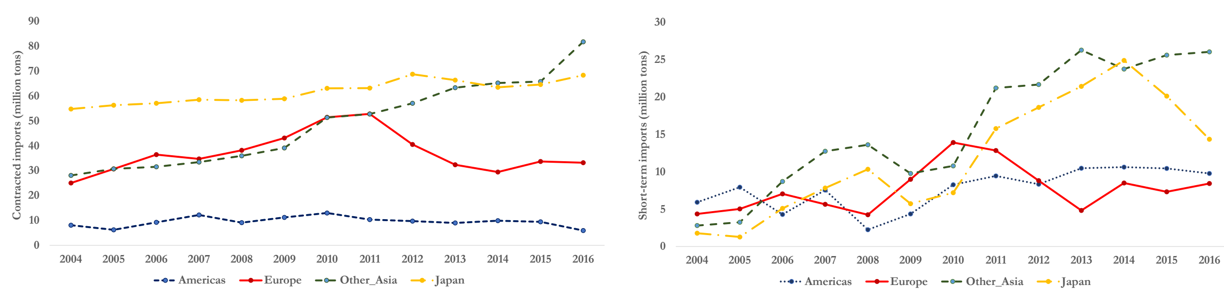

Source: GIIGNL.

A.2 Additional Descriptive Evidence

Availability of alternative buyers: This section provides additional evidence that sellers signing long-term contracts are limited in the set of buyers they can feasibly contract with, meaning that buyers may have significant bargaining power. This may seem surprising at first glance: LNG is essentially a homogeneous commodity (quality differences across different suppliers are widely regarded as minimal), an LNG tanker can in principle ship LNG to any buyer in the world with access to an import terminal, and the total number of buyers who have ever signed long-term contracts equals 106. However, the set of buyers that a seller can feasibly contract with *at any given point in time*, if they are unable to reach an agreement with their preferred contracted buyer, is fairly small, since buyers only infrequently sign long-term contracts. As Table A2 shows, there are only 12 other buyers negotiating long-term contracts during the same year that an average long-term contract is being negotiated. There is also widespread heterogeneity in contract sizes within a year. Table A2 shows that if the seller restricts consideration to other buyers who sign contracts of similar size to the contract they are signing, the number of potential alternative buyers is even smaller. For instance, less than 4 buyers (on average) sign contracts that specify a total quantity not more than 50% different from the quantity specified in the contract the seller signs with their preferred buyer. This means that switching to a different buyer is likely to entail a large adjustment in the contract quantity, making it costly for the seller to switch to an alternative buyer. Of course, such considerations apply even more so to buyers (since the set of sellers who they can contract with at any given point in time is still more restricted), so that the seller too is likely to have bargaining power.

Table A2: Availability of alternative contract partners

	All agents	Restrict to agents signing similar-sized contracts		
		$< 75\%$	$< 50\%$	$< 25\%$
Avg. no. of alternative buyers	11.98	6.02	3.76	1.91
Avg. no. of alternative sellers	5.99	3.35	2.45	1.46

Note: This table reports the average number of alternative trading partners available to buyers and sellers negotiating long-term contracts, for the sample of all long-term contracts signed in 1995 or after. The first column lists the average number of other buyers and sellers signing long-term contracts during the same year. The second to fourth columns list the average number of buyers and sellers signing long-term contracts of similar total quantity during the same year: for instance, $< 50\%$ means that the contract quantity signed by the alternative trading partners is within +/- 50% of the contract we see in the data.

Geography and ex-ante/ex-post contracting: Table 2 of Section 3 provides evidence that when sellers are far away from buyers *relative* to the buyer they are signing the contract with, they negotiate larger ex-ante contracts. The *relative distance* of each agent to their alternative trading partners, in that regression, is defined as the distance from the agent to its alternative trading partners (using the 25th percentile of the distance between that agent and all trading partners), divided by the distance to the trading partner that they are negotiating the contract with. An example may be helpful to explain how this measure is constructed. Consider the contract negotiations between Algeria's Sonatrach (an LNG exporter) and Spain's Iberdrola (an LNG importer) in 2005. The distance between the two contracting parties is 362 nautical miles. The 25th percentile of the distance between Algeria and all other buyers is 1452 nautical miles, so the relative distance for Algeria (the seller) is 1452/362 = 4.01. The 25th percentile of the distance between Spain and all other sellers is 3441 nautical miles, so the relative distance for Spain (the buyer) is 3441/362 = 9.51. In this example, the relative distance measure is large for the buyer, who is located rather far from alternative sellers, but is low for the seller, who is rather close to alternative buyers (e.g., in Europe).

Table A3 repeats the contract quantity regression shown in Table 2, but using an alternative measure of relative distance. I measure how far an agent is from alternative trading partners by the mean of the distance from all potential trading partners (instead of taking the 25th percentile of the distance, as in the original regression in Table 2). This is a reasonable measure of distance to the next best alternative partner if, in the event of a contractual breakdown, an agent was equally likely to trade with each one of the alternative partners (instead of preferring to trade with those that are nearer). I then compute the relative distance of each agent to their alternative trading partners, as the ratio of the distance from the agent to its alternative trading partners (as defined above), to the distance to the trading partner that they are negotiating the contract with. As Table A3 shows, the results remain similar: the higher the relative distance of the seller from their potential alternative buyers, the larger the size of ex-ante contracts; but the higher the relative distance of the buyer from their potential alternative sellers, the smaller the size of ex-ante contracts. This is consistent with the theoretical prediction that as sellers' outside options get worse and as buyers' outside options get better, the stronger the risk of under-investment and therefore the greater the reliance on ex-ante contracts to mitigate under-investment.

Table A4 repeats the contract quantity regression in Table 2, but with the inclusion of additional controls. Column (1) is identical to the second column of Table 2. Column (2) controls for the capacity of the export project,. The coefficient on capacity is positive (though only significant at the 10% level), which suggests that sellers building larger export projects tend to sign larger contracts with each buyer. Column (3) controls for whether or not the seller and buyer contracted in the past, finding a negative effect: this suggests potentially a desire for buyers and sellers to avoid becoming too reliant on one trading partner. In addition, I also control for the "rule of law" in both the export and import countries, a measure of judicial quality and contract enforcement developed by Kaufmann et al. (2004); Nunn (2007) had found that countries with better contract enforcement (as captured by the "rule of law") specialize more in the production of inputs requiring relationship-

Table A3: Contract quantity regressions: alternative measure of relative distance

	(1)	(2)
Dependent variable:	ln(Quantity)	
Ex-ante contract	0.69***	0.74***
	(0.14)	(0.24)
Distance	0.10***	0.091**
	(0.038)	(0.038)
Relative distance, seller	0.41***	0.31**
	(0.13)	(0.15)
Relative distance from buyer	-0.26**	-0.17
	(0.11)	(0.12)
Ex-ante*Relative distance, seller		0.55*
		(0.30)
Ex-ante*Relative distance, buyer		-0.57**
		(0.27)
Extension	-0.30*	-0.30*
	(0.16)	(0.16)
Time Trend	-0.024***	-0.024***
	(0.0055)	(0.0055)
Constant	48.7***	49.1***
	(11.0)	(11.0)
N	337	337
R^2	0.17	0.18

Note: Each observation is a long-term contract. The contract quantity is the lifetime quantity to be traded between the buyer and the seller, which is the product of the annual quantity and the contract duration. The sample includes every long-term contract that specifies a fixed export and import location (contracts with “flexible" origins or destinations are excluded). Extensions are renewals of existing contracts. Statistical significance at the 10%, 5%, and 1% levels are denoted with *, **, and ***, respectively.

specific investments. In my context, the rule of law for importers is found to negatively predict the contract quantity. In Column (4), I consider a different way of measuring the strength of the outside option for sellers and buyers, based on the availability of other buyers and sellers who sign contracts in the same year. If there are several other buyers signing contracts during the same year, this should enhance the seller’s outside option; likewise if there are several other sellers signing contracts during the same year, this should enhance the buyer’s outside option. As Column (4) shows, as the number of available buyers increases, the contract quantity decreases, consistent with the theoretical prediction that sellers sign smaller contracts when their outside option is stronger. However, the coefficient on the number of available sellers is insignificant. Finally, Column (5) includes importer and exporter region fixed effects. Across all these specifications, the effect of geography on contracting behavior is similar: ex-ante contract quantities are larger as the relative distance of the seller from alternative buyers increases and as the relative distance of the buyer from alternative sellers decreases, consistent with the predictions of the model.

I also investigate the relationship between the strength of outside options and ex-ante contracting at the project level, instead of at the contract level as in the preceding regressions. For each export project, I compute the share of capacity committed under ex-ante contracts, and regress it on the distance from the seller to nearby buyers (which is the 25th percentile of the distance from the seller to all buyers). As Table A5 shows, the coefficient on distance is positive and significant and is robust to controlling for the capacity of the liquefaction project and the overall share of capacity that is contracted (both ex-ante and ex-post). The coefficient remains positive and similar in magnitude if we control for regional dummies, but becomes insignificant. This is likely because of limited power in these regressions at the project level, since we only exploit variation across sellers (and not variation across buyers). The size of the coefficient on distance suggests that increasing the distance between an export project and nearby buyers by 1000 nautical miles will increase the share of contract quantity signed ex-ante by 8-13 percentage points.

B Estimation Details

B.1 Algorithm for solving spot market equilibrium

The spot market equilibrium in period t is characterized by a set of spot quantity choices by each seller i, $\{S_{ijt}\}_j^J$, such that the first-order condition (B1) is satisfied:

$$\underbrace{p^*_{jt} + S_{ijt}\frac{\partial p^*_{jt}(S_{ijt}, S_{-ijt})}{\partial S_{ijt}}}_{\text{Marginal revenue of selling to market } j} - \underbrace{\left(\frac{\partial C(q_{it}, K_{it})}{\partial S_{ijt}} + c^d_{ijt}\right)}_{\text{Marginal cost of selling to market } j} \leq 0 \tag{B1}$$

with equality if $S_{ijt} > 0$.

I solve for the spot market equilibrium in any given period t using the following fixed point algorithm:

1. Start with initial guesses of S_{ijt} for every i and j.

2. Update the guesses of the spot market quantities. At each iteration l, for every seller i, use the FOC, equation (B1), to solve for $\{S^l_{ijt}\}_j^J$, taking as given the spot market quantities chosen by all other sellers in the $(l-1)$th iteration.

3. Stop iterating once $||S^l_{ijt} - S^{l-1}_{ijt}|| < tol$ for each i and j, where tol is a pre-assigned tolerance level.

Table A4: Contract quantity regressions: include various controls

	(1)	(2)	(3)	(4)	(5)
Dependent variable:	ln(Quantity)				
Ex-ante contract	0.64**	0.72***	0.65***	0.12	0.85***
	(0.26)	(0.26)	(0.25)	(0.43)	(0.27)
Distance	0.072**	0.080**	0.044	0.068*	0.063
	(0.036)	(0.036)	(0.036)	(0.039)	(0.056)
Relative distance, seller	0.25**	0.31**	0.19	0.23*	0.29*
	(0.12)	(0.12)	(0.12)	(0.12)	(0.17)
Relative distance from buyer	-0.0058	-0.050	-0.0084	-0.0030	-0.029
	(0.063)	(0.067)	(0.061)	(0.063)	(0.088)
Ex-ante*Relative distance, seller	0.92***	0.79**	0.62*	0.87***	0.55*
	(0.32)	(0.33)	(0.32)	(0.33)	(0.33)
Ex-ante*Relative distance, buyer	-0.96***	-0.83***	-0.71***	-0.89***	-0.60**
	(0.26)	(0.27)	(0.26)	(0.28)	(0.28)
log(Capacity)		0.16*			
		(0.086)			
Contracted in the past			-0.35**		
			(0.15)		
Rule of law, exporter			0.065		
			(0.066)		
Rule of law, importer			-0.44***		
			(0.12)		
No. available buyers				-0.043*	
				(0.023)	
No. available sellers				0.0047	
				(0.037)	
Ex-ante*No. available buyers				-0.0074	
				(0.039)	
Ex-ante*No. available sellers				0.070	
				(0.063)	
Extension	-0.29*	-0.25	-0.14	-0.29*	-0.27*
	(0.16)	(0.16)	(0.16)	(0.16)	(0.16)
Time Trend	-0.024***	-0.025***	-0.027***	-0.011	-0.026***
	(0.0055)	(0.0054)	(0.0055)	(0.0083)	(0.0055)
Importer region fixed effects					Yes
Exporter region fixed effects					Yes
N	337	337	337	337	337
R^2	0.19	0.20	0.25	0.22	0.23

Note: Each observation is a long-term contract. The contract quantity is the lifetime quantity to be traded between the buyer and the seller, which is the product of the annual quantity and the contract duration. The sample includes every long-term contract that specifies a fixed export and import location (contracts with “flexible" origins or destinations are excluded). Extensions are renewals of existing contracts. Statistical significance at the 10%, 5%, and 1% levels are denoted with *, **, and ***, respectively.

Table A5: Regression of share of contract volume signed before final investment decision date on characteristics of the export project

	(1)	(2)	(3)	(4)
	Dependent variable: Share of contract quantity signed before final investment decision			
Distance from nearby buyers	0.11**	0.13***	0.13***	0.082
	(0.041)	(0.041)	(0.038)	(0.049)
Capacity		-0.027**	-0.029**	-0.020
		(0.012)	(0.011)	(0.013)
Share of capacity contracted			0.43***	0.44***
			(0.13)	(0.14)
Atlantic				0.18
				(0.29)
Middle East				-0.087
				(0.24)
Pacific				0.092
				(0.23)
N	58	58	58	58
R^2	0.11	0.18	0.31	0.34

Note: Each observation is an investment project. The sample includes every investment whose final investment decision was made in 1995 or later. The distance from nearby buyers is measured in 1000 nautical miles. Statistical significance at the 10%, 5%, and 1% levels are denoted with *, **, and ***, respectively.

B.2 Demand Estimates: Additional Details and Results

This section provides further details of demand estimation, as well as some robustness checks. Table B6 shows the baseline demand estimates.

In the baseline demand specification (Table B6), I assume that Q_{jt}/R_{jt} is linear in the spot price, where R_{jt} is the maximum regasification capacity ever reached by country j. The most natural measure of R_{jt} is the regasification capacity operational in country j in period t, since this directly measures the physical import capacity of the country. However, using regasification capacity directly in the denominator of the LHS of equation (10) has the disadvantage of assuming that a country's demand for LNG increases the instant it builds new capacity, which is unlikely to be realistic and would imply sharp swings in demand profiles for countries over time. Instead, R_{jt} is assumed to equal the maximum regasification capacity built by country j during the sample period. While this may appear to be an unusual modelling choice, the demand estimates are quite similar under alternative specifications of the dependent variable, as illustrated in Table B7.

Column (1) is the baseline specification. In Column (2), R_{jt} is the regasification capacity of country j at time t. In Column (3), R_{jt} is the maximum LNG import of country j over the sample

Table B6: Demand Curve Estimates

	1st-stage		2SLS	
	Estimate	S.E.	Estimate	S.E.
Price			-0.029**	(0.014)
log(Elec. Cons., fossil)	0.15	(0.68)	0.31***	(0.10)
Min. temp	0.0081	(0.014)	-0.0042**	(0.0021)
Oil Price	0.11***	(0.0080)	0.0031**	(0.0015)
Excluded IVs				
Elec cons. else	0.0061***	(0.00061)		
Min. temp else.	-0.51***	(0.054)		
N	815		815	
Kleibergen-Paap F-stat	61.0	Olea-Pflueger Effective F-stat		35.6
Mean elasticity	-1.77	Median elasticity		-0.92

Note: Each observation is an importer-year-quarter pair. The second and third columns report the first-stage regression of the spot price on the controls and instruments. The last two columns report 2SLS estimates of the demand curve, where the dependent variable in is total LNG imports in country j in period t divided by the maximum regasification capacity reached by country j between 2004 and 2017. All regressions include importing country fixed effects. The instruments for prices are total electricity consumption from fossil fuels in period t excluding country j's own consumption (“Elec cons. else"), as well as the average of the minimum temperature in period t for all importing countries excluding country j (“Min temp else"). Effective F-statistics for excluded IVs following Olea and Pflueger (2013) are reported, as well as the Kleibergen-Paap F-statistic. Standard errors are clustered by country.

period. The demand estimates in columns (2) and (3) are qualitatively similar to that in column (1). Moreover, the estimated mean and median demand elasticities are similar in magnitude in the three specifications.

A potential drawback of the demand system in (10) is that it does not explicitly take into account capacity constraints faced by the buyer. Equation (10) in principle allows an importer to import more LNG than they have the capacity to process, whereas under normal circumstances, the regasification capacity of a country provides an upper limit to how much LNG the country can import.[45] These capacity constraints are very rarely close to binding: import capacity utilization is typically quite low (41% on average), and is less than 90% for more than 94% of the observations. Still, one might wonder if the demand estimates are significantly affected by the few observations where capacity utilization is high and buyers may be constrained in their ability to import additional LNG.

As such, I also estimate a version of the demand function which explicitly constrains buyers' capacity utilization to be less than 1. Let $s_{jt} = Q_{jt}/R_{jt}$ denote buyer j's capacity utilization in period t, where with some abuse of notation I now use R_{jt} to denote buyer j's regasification capacity in period t. I define the dependent variable in the demand equation to be a logit transformation of s_{jt},

[45] Under exceptional circumstances, countries can import LNG in excess of their capacity, as was the case for Taiwan from 2012 to 2017.

or $ln(s_{jt}/(1-s_{jt})$, as in equation (B2). The logit transformation ensures that capacity utilization, or s_{jt}, must lie strictly between 0 and 1.

$$ln\left(s_{jt}/(1-s_{jt})\right) = \alpha - bp_{jt} + \theta_j + x_{jt}\theta_{dx} + \varepsilon_{jt} \tag{B2}$$

The demand estimates from this "logit" specification are shown in the fourth column of Table B7. All of the coefficients have the same sign as in Column (1). Moreover, the estimated demand elasticities are not too different: demand is somewhat more elastic for the median observation (-1.21 rather than -0.92) and somewhat less elastic on average (-1.36 rather than -1.77). Because it is computationally much more challenging to solve for the spot market equilibrium when demand is non-linear, the specification that I use for the majority of the analysis is the simpler linear specification from Column (1).

Table B7: Demand Estimates: Different Dependent Variables

	(1)	(2)	(3)	(4)
Dependent variable	Imports/Max Regas	Imports/Regas	Imports/Max Imports	Logit
Price	-0.029**	-0.052***	-0.023**	-0.26***
	(0.014)	(0.016)	(0.0090)	(0.078)
log(Elec. Cons., fossil)	0.31***	0.35***	0.24***	1.60***
	(0.10)	(0.075)	(0.077)	(0.41)
Min. temp	-0.0042**	-0.0066***	-0.0036***	-0.031***
	(0.0021)	(0.0024)	(0.0013)	(0.012)
Oil Price	0.0031**	0.0056***	0.0024***	0.028***
	(0.0015)	(0.0017)	(0.00093)	(0.0082)
N	815	815	815	815
R^2	0.40	0.26	0.35	0.28
Kleibergen-Paap F-stat	61.0	61.0	61.0	61.0
Effective F-stat	35.6	35.6	35.6	35.6
Mean elasticity	-1.77	-1.82	-1.71	-1.36
Median elasticity	-0.92	-1.20	-0.87	-1.21
Other Controls	Importer fixed effects			

Note: Each observation is an importer-year-quarter pair. The control variables, fixed effects and instruments are the same as in Table B6. Standard errors are clustered by country.
The dependent variable in column (1) is total LNG imports in country j in year-quarter t, measured in million tonnes, divided by the maximum regasification capacity reached by country j between 2004 and 2017. In column (2), the dependent variable is LNG imports divided by the regasification capacity of country j in period t, while in column (3) it equals LNG imports in period t divided by the maximum LNG imports by country j in the sample. Finally, column (4) presents the logit specification from equation (B2).

Next, I study how the choice of instruments affects the demand estimates. In addition to the two instruments I use for the baseline demand estimation (minimum temperature and electricity consumption from fossil fuels in rival countries), I consider two additional instruments. The first of these is total electricity generation from non-fossil fuel sources (nuclear, renewables, hydro) in rival countries (i.e., all importers apart from j). When electricity generation from these sources increases, then rival countries are likely to reduce their demand for LNG, which in turn reduces the spot price paid by country j. For this instrument to be valid, it has to be the case that changes in electricity generation from non-fossil fuel sources in rival countries is uncorrelated with idiosyncratic LNG demand shocks in country j.

Table B8: Demand Estimates: Different IVs

	(1)	(2)	(3)	(4)	(5)	(6)
Price	-0.029**	-0.024**	-0.031***	-0.034***	-0.026**	-0.031***
	(0.014)	(0.011)	(0.010)	(0.011)	(0.012)	(0.011)
log(Elec. Cons., fossil)	0.31***	0.31***	0.31***	0.31***	0.31***	0.31***
	(0.10)	(0.099)	(0.10)	(0.10)	(0.100)	(0.10)
Min. temp	-0.0042**	-0.0038**	-0.0043**	-0.0046**	-0.0040**	-0.0044**
	(0.0021)	(0.0018)	(0.0018)	(0.0019)	(0.0018)	(0.0018)
Oil Price	0.0031**	0.0026**	0.0033***	0.0037***	0.0029**	0.0034***
	(0.0015)	(0.0012)	(0.0011)	(0.0012)	(0.0012)	(0.0011)
N	815	815	815	815	815	815
R^2	0.40	0.44	0.38	0.35	0.42	0.38
Kleibergen-Paap F-stat	61.0	52.1	75.0	97.2	52.9	65.1
Effective F-stat	35.6	46.1	38.1	41.8	39.6	31.4
Other Controls	Importer fixed effects					
Instruments						
Elec. cons. else	Y	Y	Y	Y	Y	Y
Min. temp else.	Y		Y	Y		Y
Elec gen. else				Y	Y	Y
Global Liq. Cap.		Y	Y		Y	Y

Note: Each observation is an importer-year-quarter pair. The spot price is measured in \$/MMBtu. The dependent variable is total LNG imports in country j in year-quarter t, measured in million tonnes, divided by the maximum regasification capacity reached by country j between 2004 and 2017. Effective F-statistics for excluded IVs following Olea and Pflueger (2013) are reported, as well as the Kleibergen-Paap F-statistic. Standard errors are clustered by country.
The four instruments for prices are (i) total electricity consumption from fossil fuels in period t excluding country j's own consumption ("Elec cons. else") (ii) the average of the minimum temperature in period t for all importing countries excluding country j ("Min temp else") (iii) total electricity generation from non-fossil fuel sources (nuclear, renewables, hydro) in period t excluding country j ("Elec gen. else") (iv) total liquefaction capacity in the world in period t ("Global Liq. Cap."). Each column presents estimates from a different combination of instruments.

The second is total liquefaction capacity (i.e. total export capacity) in the world in period t (which is abbreviated to "Global Liq. Cap." in the tables). 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 from fossil fuels. 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.

Table B8 presents demand estimates from different combinations of these four instruments. I am unable to obtain a strong first-stage without including electricity consumption from fossil fuels in rival countries as an instrument, so this instrument appears to be particularly pivotal and I include it in all of the specifications. Varying the remaining instruments makes little difference to the estimated coefficients, which is evident from comparing the different columns in Table B8. Finally, the demand estimates are also robust to alternative measures of weather and different choices of controls and fixed effects.[46]

B.3 Cost Estimates

Table B9 presents cost function estimates, which were discussed in Section 5.2.

Table B9: Cost Parameter Estimates

Cost Parameters	Estimate	S.E.	
δ	17.32	(0.13)	
ν (calibrated)	*2*		
Fit (R^2)			
Prices p_j	0.57	Production q_i	0.93
Spot Trade Flows S_{ij}	0.13	Regional Spot Trade Flows	0.47

Note: Each observation is an exporter-importer-year pair (N = 3245). Regional spot trade flows are the spot trade flows aggregated to the regional level, with importers and exporters divided into 9 separate regions (e.g., Northeast Asia, Southeast Asia etc.). Heteroskedasticity-robust standard errors reported in parentheses.

While the model is parsimonious, the fit of the model is reasonably good, as shown in the last panel of Table B9. The R^2 for prices is 0.57. Appendix Figure B3 shows actual versus predicted prices in Northeast Asia and Southwest Europe, the two main LNG-importing regions in the world. We can see that the model generates a reasonable fit for year-to-year changes in prices. The R^2 for

[46]These results are omitted for brevity, but available upon request.

total exports is 0.93, though this is partly because total exports are to a large extent determined by long-term contracts. The R^2 for spot trade flows is on the low side at 0.13, but the model provides a good fit ($R^2 = 0.47$) for spot trade flows aggregated to the regional level (e.g., spot exports from Middle East to Northeast Asia). This is because buyers in the same region are often located close to one another (e.g., Japan and South Korea), and such buyers will end up paying almost identical spot prices (both in reality and according to the model). The model finds it difficult to predict the exact allocation of spot LNG between buyers in the same region, since sellers should be nearly indifferent between selling to such buyers. However, the model is able to explain well how sellers allocate LNG across geographically segmented regions.

Figure B3 shows how the model-predicted spot prices compare with actual spot prices.

Figure B3: Actual vs. Predicted Spot Prices

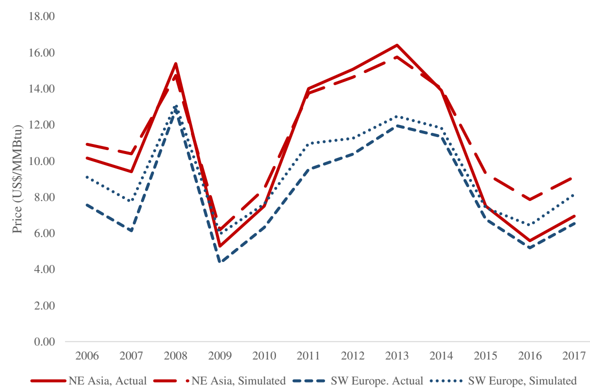

Note: The figure plots actual spot prices as well as model-predicted annual spot prices in North-east Asia (China, Japan, South Korea, Taiwan) and Southwest Europe (France, Greece, Italy, Spain). The spot price in each region is calculated as a quantity-weighted average of the spot price paid by each individual country in the region.

In the baseline analysis, I assumed firms face soft capacity constraints, with the marginal cost is quadratic in capacity utilization q/K (ν=2 in equation (11). Here I explore two alternative specifications of the cost function. First, I look at the effect of assuming $\nu = 1$ (column (1) of Table B10), meaning that marginal cost is linear (rather than quadratic) in capacity utilization. These results are similar to the baseline model (Table B9), with the estimated δ of 15 implying a marginal cost of $15/MMBtu at 100% capacity utilization. However, the model fit is a little worse, since unlike the baseline model, this model does not penalize excess capacity utilization very much, and so firms end up adjusting their production more flexibly than we can see in the data.

Second, I consider a different way to model capacity constraints that has been popular in the empirical literature. I assume each firm faces a “hockey-stick" cost function, with its marginal costs constant until its capacity utilization hits a threshold of $\tilde{\nu}$. If the firm's capacity utilization exceeds $\tilde{\nu}$, the firm's marginal costs increase with the level of capacity utilization in excess of the

Table B10: Cost Parameter Estimates

	(1)	(2)
	MC linear in capacity utilization	Hockey-stick
δ	15.06	738.27
	(0.09)	(15.72)
ν	1	
Hockey-stick threshold, $\tilde{\nu}$		0.90
N	3245	3245
Buyers	40	40
Fit (R^2)		
Prices p_j	0.45	0.71
Total Exports by Exporter q_i	0.93	0.95
Spot Trade Flows S_{ij}	0.11	0.29
Regional Spot Trade Flows S_{ij}	0.39	0.67

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. Heteroskedasticity-robust standard errors reported in parentheses. In Column (1), the cost function follows equation (11), with the parameter ν calibrated to 1. In Column (2), the cost function is of the hockey-stick form (equation (B3)), with the "threshold" parameter ν calibrated to 0.90.

threshold, with the parameter δ governing the rate at which marginal costs increase as the firm approaches full capacity utilization. This follows Ryan (2012) and Miller and Osborne (2014), who use a similar specification of the cost function in their analyses of the US Portland cement industry.

$$C(q_{it}, K_{it}) = \frac{\delta}{2} 1(q_{it} \geq \tilde{\nu} K_{it})(q_{it} - \tilde{\nu} K_{it})^2 \tag{B3}$$

I estimate δ in the above equation, fixing ν at 0.90 (meaning that costs start to increase at 90% capacity utilization). As shown by column (2) of Table B10, this model leads to a better fit of prices and spot trade flows than the baseline model. However, these estimates imply economically implausible costs of high capacity utilization: for the average seller, exceeding the 90% capacity utilization threshold even slightly leads to very high marginal costs that exceed the maximum prices ever observed in the data, whereas in practice we sometimes see sellers operating at or near 100% capacity utilization. As a consequence, I adopt the model of "soft" capacity constraints (as described in the main text) which leads to more plausible estimates of the cost of operating at high levels of capacity utilization.

B.4 Estimates of contracting and investment parameters: further details

This section presents additional estimates of the parameters characterizing contracting and investment behavior. In Table B11, I report the results from specifications where the Nash bargaining weight τ is allowed to vary across seller and buyer groups.[47] In specification (1), I allow sellers that are national oil and gas companies (NOCs) in the 4 traditional major LNG exporting countries (Qatar, Algeria, Indonesia and Malaysia) to have a different bargaining parameter from all other sellers. I find little difference in the bargaining power of these two groups of sellers; one reason is that international oil companies are heavily involved in LNG exporting, and so even the "smaller" sellers may in practice have considerable bargaining leverage. In specification (2), I allow sellers to have a different bargaining weight when negotiating with North-east Asian buyers (Japan, China, South Korea, Taiwan) than with other importers. Unlike LNG buyers in Europe, these buyers have limited ability to switch between LNG and piped natural gas (with the exception of China), which may lessen their bargaining power in a way that is not directly captured in the model. Consistent with that, I find that the seller's Nash bargaining weight τ is estimated to be somewhat higher when negotiating with North-east Asian buyers than with other buyers, though the difference is not statistically significant. Finally, specification (3) allows τ to differ across both seller and buyer groups, with much the same results as in the first two specifications. Overall, these estimates suggest that differences in bargaining power across sellers and buyers, though they may exist, do not appear to be sizeable and are difficult to estimate with precision. As such, I use the baseline specification with a single bargaining weight τ for the main analysis.

Next, I investigate determinants of the "contract premium" (i.e., buyer's additional WTP to purchase LNG under long-term contracts as opposed to buying on the spot market). As Column (1) of Table B12 shows, on average there is evidence of a contract premium in Stage 3 (for ex-post contracts), but not in Stage 1 (ex-ante). In column (2), we allow the contract premium to differ for buyers located in the Asia-Pacific region: unlike buyers in Europe and North America, these buyers tend to have less access to pipeline gas, so they may in principle have different preferences for long-term contracts; however, I find no significant difference in the contract premium they are willing to pay compared to the contract premium of other buyers.

Column (3) includes the "rule of law" in both the export and import countries, a measure of judicial quality and contract enforcement developed by Kaufmann et al. (2004). This is in part motivated by Nunn (2007)'s finding that countries with better contract enforcement (as captured by the rule of law) specialize more in the production of inputs requiring relationship-specific investments. In a similar vein, one might expect that superior rule of law might lead to a preference for signing long-term contracts, since there is a reduced probability of ex-post opportunism and/or

[47] Since each agent (whether a seller or a buyer) signs a relatively small number of contracts, it is not feasible to precisely estimate bargaining parameters separately for every agent.

Table B11: Contracting and investment parameter estimates: determinants of bargaining power

	Spec (1)		Spec (2)		Spec (3)	
	Estimate	S.E.	Estimate	S.E.	Estimate	S.E.
Investment cost, γ_1	89.76	(3.55)	89.59	(3.66)	89.75	(3.74)
Investment cost, γ_2	-0.31	(0.32)	-0.28	(0.40)	-0.31	(0.37)
Contract premium (ex-ante) : κ_1	-0.008	(0.010)	-0.008	(0.011)	-0.008	(0.011)
Contract premium (ex-ante) : θ_1	0.18	(0.46)	0.18	(0.46)	0.13	(0.46)
Contract premium (ex-post): κ_3	-0.009	(0.006)	-0.009	(0.006)	-0.009	(0.006)
Contract premium (ex-post): θ_3	1.40	(0.12)	1.40	(0.12)	1.40	(0.12)
Bargaining weight, τ:						
τ(major NOC exporters)	0.62	(0.08)				
τ(other exporters)	0.69	(0.06)				
τ(other importers)			0.44	(0.57)	0.45	(1.17)
τ(NE-Asian importers)			0.64	(0.22)		
τ(NE-Asian importers*major NOC exporters)					0.62	(0.30)
τ(NE-Asian importers*other exporters)					0.69	(0.11)
Number of contracts (Stage 3)	172		172		172	
Number of contracts (Stage 1)	123		123		123	
Number of investments	54		54		54	

Note: In these specifications, τ is allowed to differ across seller and buyer groups. “Major NOCs" refers to Qatar, Algeria, Indonesia and Malaysia (the 4 largest exporters that have national oil companies), while “other exporters" refers to all other exporters. “NE-Asian importers" includes Japan, China, South Korea and Taiwan, while “other importers" refers to all other importers. All parameters estimated using non-linear least squares. Standard errors are heteroskedasticity-robust.

breach. Consistent with that, I find that there is a positive relationship between the ex-post contract premium (in Stage 3) and the rule of law in importers and exporters alike: the coefficient on the rule of law in exporting countries is significant at a 5% level, whereas the coefficient on the rule of law in importing countries is significant at a 10% level. There is little evidence, however, that the rule of law affects ex-ante contracting decisions. Finally, column (4) includes an indicator for whether the buyer and seller contracted in the past, finding that it makes little difference to the contract premium.

C Counter-factual Details

C.1 Implementation of Counter-factual Analyses

I assume across all counter-factual experiments that in a given investment project, the seller *i* contracts with at most one buyer (as discussed in Section 6. For some projects, the seller is observed in the data to negotiate contracts with multiple buyers, some in Stage 1 and Stage 3. While multiple

Table B12: Contracting and investment parameter estimates: determinants of contract premium

	Spec (1)		Spec (2)		Spec (3)		Spec (4)	
	Estimate	S.E.	Estimate	S.E.	Estimate	S.E.	Estimate	S.E.
Investment cost, γ_1	89.76	(3.55)	89.96	(3.64)	89.99	(3.61)	90.05	(3.58)
Investment cost, γ_2	-0.31	(0.32)	-0.34	(0.29)	-0.34	(0.29)	-0.35	(0.28)
Contract premium (ex-ante)								
κ_1	-0.008	(0.010)	-0.006	(0.011)	-0.007	(0.011)	-0.010	(0.012)
θ_1	0.18	(0.46)	0.44	(0.62)	0.76	(1.25)	0.88	(1.31)
θ_1: Importer, Pacific			-0.51	(0.67)	-0.45	(0.87)	-0.25	(0.97)
θ_1: rule of law, importer					-0.23	(0.47)	-0.14	(0.49)
θ_1: rule of law, exporter					-0.10	(0.22)	-0.09	(0.22)
θ_1: contracted in the past							-0.48	(0.60)
Contract premium (ex-post)								
κ_3	-0.009	(0.006)	-0.011	(0.006)	-0.003	(0.006)	-0.005	(0.006)
θ_3	1.40	(0.12)	1.02	(0.34)	0.28	(0.73)	0.47	(0.93)
θ_3: Importer, Pacific			0.51	(0.35)	0.34	(0.35)	0.39	(0.37)
θ_3: rule of law, importer					0.60	(0.34)	0.67	(0.38)
θ_3: rule of law, exporter					0.27	(0.12)	0.28	(0.12)
θ_3: contracted in the past							-0.35	(0.64)
Bargaining weight, τ:								
τ: major NOC exporters	0.62	(0.08)	0.62	(0.07)	0.62	(0.07)	0.62	(0.07)
τ: other exporters	0.69	(0.06)	0.70	(0.06)	0.70	(0.07)	0.70	(0.07)
No. of contracts (Stage 3)	172		172		172		172	
No. of contracts (Stage 1)	123		123		123		123	
No. of investments	54		54		54		54	

Note: “Major NOCs" refers to Qatar, Algeria, Indonesia and Malaysia: the 4 largest exporters that have national oil companies. “North-east Asian importers" includes Japan, China, South Korea and Taiwan, while “other importers" refers to all other importers. In Spec. 2, “Importer, Pacific" is an indicator for buyers in the Asia-Pacific region (that is, buyers located in Asia, Middle East or South America). All parameters estimated using non-linear least squares. Standard errors are heteroskedasticity-robust.

buyers are accommodated in the Nash-in-Nash bargaining model and pose no difficulty for estimation, they create computational difficulties in the counter-factual analysis: solving for optimal contract quantities and investments (which needs to be done numerically) turns out to be computationally intensive when there are multiple buyers. As such, for any projects with more than one buyer, I select the buyer with the largest contract quantity, and assume that the seller can only negotiate with the buyer in the counter-factual simulations. Reassuringly, the baseline simulations under this assumption yields similar investment levels to what I see in the data. Partly this is because even though the seller is only permitted to sign a contract with one buyer, the seller can still sign a large contract with that one buyer (if they wish to), and so this restriction does not much affect the seller’s incentive to invest.

In the partial equilibrium counter-factual analyses of Section 6.1, I solve for sub-game perfect equilibrium contracting and investment choices for each investment project (consisting of a single seller i and a single buyer j), holding fixed their beliefs Y_{-i} about the contracting and investment choices of the rest of the sellers and buyers. To find the equilibrium of the multi-stage game, I search numerically for the investment K_i and the contract quantities $q_{ij}^{c,1}$ and $q_{ij}^{c,3}$ that satisfy the first-order conditions (5), (8), and (9). In counter-factuals with no contracting permitted, this is simpler as I only need to solve for the investment level K_i such that the investment first-order condition, equation (8), is satisfied.

In the general equilibrium counter-factuals described in Section 6.3 and Section 6.4, I solve for the full industry equilibrium, through the following fixed point algorithm:

1. Start with initial guesses of investment K_i, the contract quantity in Stage 1, $q_{ij}^{c,1}$, and the contract quantity in Stage 3, $q_{ij}^{c,3}$, for every seller i and buyer j considered in the counter-factual.

2. Update the guesses of the investment and contract quantities. At each iteration l, for every seller i and their contracted buyer j:

    - Update the beliefs of seller i and buyer j about investment and contracting by the rest of the industry to $Y_{-i}^{l-1} = \{q_{-i}^{c,l-1}, K_{-i}^{l-1}\}$, using the investment and contracting choices from the previous iteration, or iteration $l-1$.

    - Solve numerically for new guesses of investment K_i^l, the Stage 1 contract quantity $q_{ij}^{c,1,l}$ and the Stage 3 contract quantity $q_{ij}^{c,3,l}$ that satisfy the first-order conditions (5), (8), and (9) for seller i and buyer j, using their new beliefs Y_{-i}^{l-1}. (The superscript l refers to the the fact that these are guesses for the $l-th$ iteration).

3. Stop iterating once $||K_i^l - K_i^{l-1}|| < tol$, $||q_{ij}^{c,1,l} - q_{ij}^{c,1,l-1}|| < tol$, and $||q_{ij}^{c,3,l} - q_{ij}^{c,3,l-1}|| < tol$ where tol is a pre-assigned tolerance level.

C.2 Contracting and allocative efficiency

Section 6.2 showed that allocative efficiency is higher when no long-term contracts are used compared to the baseline, where there are no restrictions on long-term contracting. This section further explores the allocative efficiency effects of using long-term contracts, by investigating how welfare varies with the extent of long-term contracting. I extend the analysis in Section 6.2 by considering intermediate regimes where there is neither unrestricted long-term contracting, nor a complete ban on long-term contracting.

The motivation behind this analysis is that the allocative efficiency effects of long-term contracts depend on the extent of seller contractual commitments (see Section 4.4, page 25). Long-term contracts reduce the flexibility of sellers in meeting demand shocks, but decrease distortions from market power in the spot market. If contractual commitments account for a large share of seller capacity, there is very little spare capacity on the spot market that can be deployed to flexibly deal with demand fluctuations, so the former effect dominates and the allocative efficiency gains from reducing contract usage are more likely to be large. As the share of capacity that is contracted decreases, though, the flexibility gain from freeing up capacity is likely to be smaller, so that at some point we would expect the market power effect to dominate. Thus it seems unlikely that allocative efficiency is maximized by eliminating long-term contracts entirely. The purpose of this section is to explore the allocative efficiency effects of using different mixes of contracts and spot trade.

As such, in addition to the two extreme regimes considered in Section 6.2 (no contracting, unrestricted contracting), I also study intermediate regimes where contracting is permitted, but the contract quantity is reduced relative to the baseline. I consider three intermediate regimes, in each of which the contract quantity is reduced to X% of the originally agreed to quantity (where X can be 75, 50 or 25). For example, in the “75% contracts regime", the quantity agreed to in each contract is reduced by 25%, so that a 4 mtpa contract is transformed into a 3mtpa contract.[48]

Table C13: Allocative efficiency with different levels of long-term contracting, holding investment fixed

	Benchmark	75% contracts	50% contracts	25% contracts	No contracts
Welfare ($ bn)	5,469	5,496	5,508	5,511	5,507
Seller surplus ($ bn)	2,535	2,557	2,572	2,583	2,517
Buyer surplus ($ bn)	2,934	2,939	2,937	2,928	2,990

Note: In the benchmark regime, sellers and buyers can sign long-term contracts in either Stage 1 or Stage 3. In every other regime, we reduce the size of each contract relative to the baseline. For example, in “75% contracts", the quantity agreed to in each contract is reduced by 25% (so as to equal 75% of the baseline quantity). In the “no-contracts" regime, no contracting is permitted. Investment is held fixed (at observed levels) across all the counter-factuals. The table shows discounted total welfare, total seller surplus and total buyer surplus in US$bn from 2001 onwards.

As Table C13 shows, reducing contract usage from the baseline (no restrictions) to the “75% contracts" leads to a large welfare gain, of $27 bn. Diminishing returns quickly kick in, however, so that going from “75% contracts" to “50% contracts" improves welfare by $12 bn, and going from

[48] It is important to note that these counter-factuals are somewhat artificial in nature and do not provide a prescription for actual policy: it is hard to imagine how a policymaker would impose a 25% or 50% reduction in contract usage. The point of these counter-factuals is simply to figure out if intermediate levels of contracting can be used to achieve higher allocative efficiency.

“50% contracts" to “25% contracts" only improves welfare by $3bn. Beyond that point, reducing contract usage (from “50% contracts" to “25% contracts") in fact reduces welfare by $4bn. Figure C4 graphically illustrates this non-monotonic relationship between the level of contracting and allocative efficiency.

Figure C4: Contracts and Allocative Efficiency

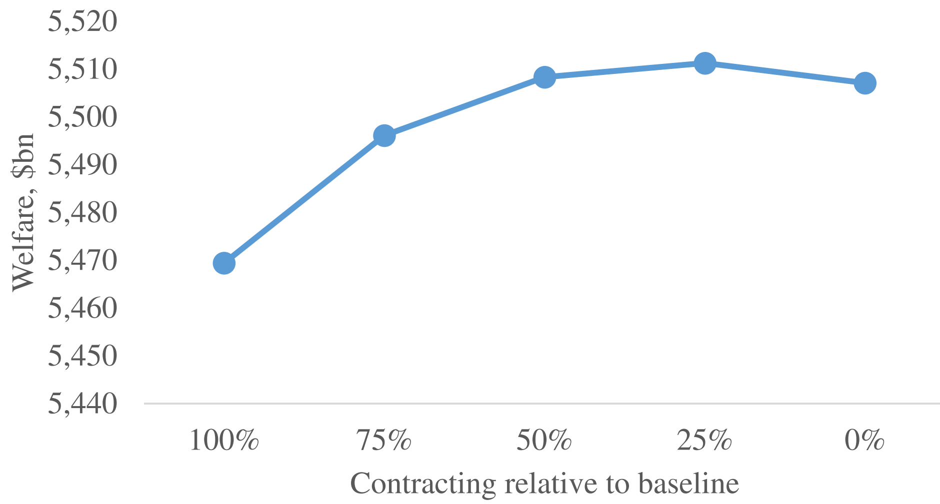

Note: The vertical axis plots industry welfare (in US$ bn). The horizontal axis shows the level of contracting relative to the baseline: for instance, 75% means that the quantity agreed in every contract is shrunk to 75% of the baseline contract quantity (so that a 4mtpa contract becomes a 3mtpa contract). See Table C13 for a more detailed description.

C.3 Industry responses to demand shocks

In this section, I describe two counter-factual exercises designed to explore the efficiency of the LNG industry's responses to demand shocks with and without long-term contracts. My goal is to investigate whether the use of long-term contracts limits the ability of sellers to respond efficiently to demand shocks in the short run.

Fukushima nuclear disaster The first of these counter-factuals explores the effect of the Fukushima nuclear disaster in Japan (in March 2011) on the LNG industry. In the months following the Fukushimna disaster, Japan shut down all of their nuclear plants, which had previously accounted for around 25% of its total electricity generation capacity.[49] The resulting shortfall was met by a combination of increased imports of fossil fuels (most notably LNG and oil), as well as demand conservation measures designed to lessen the use of electricity (Miyamoto et al., 2012; Neidell

[49]Calculations based on 2010 data, which was obtained from BP's Statistical Review of World Energy, 2021.

et al., 2019). In order to model the effect of this disaster on the LNG market, I assume that the entirety of the increase in Japan's electricity generation from natural gas in 2011 - 2013 (relative to 2010) was a consequence of the Fukushima disaster. This is a reasonable assumption as electricity generation from natural gas had been very stable between 2008 and 2010, the years preceding the Fukushima disaster.

I consider a number of alternative contracting regimes. The baseline takes the industry as it was, with no restrictions on long-term contracting. The "No Contracts" regime assumes that no long-term contracts are binding, so that all of trade is carried out via the spot market. "No Japanese Contracts" assumes that only the long-term contracts agreed to by Japanese buyers are non-binding, but all other contracts are binding. "No non-Japanese Contracts" reverses that and assumes instead that only the long-term contracts agreed to by Japanese buyers are binding, but all other contracts are binding. Finally, I compare all of these regime with a "Fully Competitive" regimes where no long-term contracts are used and all sellers behave competitively, which provides the efficient benchmark against which we can compare all the other contracting regimes. The "Fully Competitive" regime can also be thought of as an approximation of a scenario where no destination clauses or other resale restrictions are used in long-term contracts, as discussed in Section 6.4. For each of these counter-factuals, I numerically solve for the industry equilibrium both with and without the Fukushima nuclear disaster.

Table C14: Effect of Fukushima nuclear crisis on trade and industry surplus, 2011 - 2013

	Baseline	No Contracts	No Japanese Contracts	No non-Japanese Contracts	Fully Competitive
Welfare ($bn)	507.2	511.8	507.2	511.9	522.3
Seller Surplus ($bn)	227.8	258.5	236.9	251.0	290.3
Buyer Surplus ($bn)	279.4	253.3	270.3	261.0	232.1
Japan's total LNG imports (mt)	260.2	262.7	256.8	266.8	266.2
Increase from Fukushima (mt)	49.4	51.8	45.9	56.0	55.3

Note: Welfare, seller surplus and buyer surplus are the sum of the surplus in 2011, 2012 and 2013. The fourth row reports Japan's total LNG imports, added across 2011 to 2013. The fifth row reports the increase in Japan's total LNG imports induced by Fukushima, by subtracting from the fourth row Japan's total LNG imports between 2011 and 2013 had the Fukushima crisis not taken place. I assume that the Fukushima nuclear crisis caused Japan's demand for natural gas for electricity generation to increase by 53 TWh in 2011, 109 TWh in 2012 and 98 TWh in 2013.

Table C14 illustrates the results. As we can see from comparing "Baseline" with "No Contracts", the industry responds more efficiently to the demand shock when sellers are unencumbered by long-term contracts: welfare between 2011 and 2013 is $507.2bn in the baseline regime, but rises to $511.8 bn without any long-term contracts, a $4.6bn welfare gain. This is because in the absence of long-term contracting, sellers sell more LNG to Japan: in the baseline regime, Japan's LNG imports rise by 49.4 million tonnes (mt) due to the Fukushima crisis, but would have risen even more (by 51.8 mt) if sellers were not bound by long-term contracts. The results from the "No non-Japanese Contracts" regime illustrate that these welfare gains accrue mainly from eliminating contracts signed by non-Japanese buyers, since it allows the sellers on these contracts to instead re-direct some of their LNG to Japan. By contrast, removing Japanese contracts does not lead to any efficiency gains at all; indeed, it causes Japanese LNG imports to decline relative to the baseline scenario, since sellers now have a stronger incentive to withhold LNG from Japan (due to the removal of the pro-competitive effect of contracts identified by Allaz and Vila, 1993).

Finally, the presence of market power also means that elimintating long-term contracts is not enough to achieve the efficient short-run allocation (even if it does lead to some efficiency gains). As the final column of the table shows, if LNG were competitively allocated and there were no contracts, the industry would have responded to Fukushima by selling even more LNG to Japan (a 55.3 mt increase), resulting in even higher welfare ($522 bn, or a $15 bn increase relative to the baseline regime).

Shutdown of Russian natural gas exports to Europe Second, I consider the effect of a hypothetical shutdown of Russian natural gas exports to Europe. This issue has gained considerable policy prominence in 2022, after the Russian invasion of Ukraine and the resulting geopolitical tensions. The macroeconomic consequences of a potential shutdown of European natural gas imports from Russia have been studied by Bachmann et al. (2022) and Pescatori et al. (2022). Here I focus more narrowly on the efficiency of the LNG industry response to such an event, under different contracting regimes, assuming that a Russian shutdown of natural gas exports takes place in 2017 (the last year for which I have complete data on the industry). I assume that if Europe were to stop importing natural gas from Russia, Europe's demand for LNG would rise by 100 bcm (or 73.5 mt) at baseline prices. Although Europe's total natural gas imports from Russia were considerably higher (at 155 bcm in 2021), infrastructure constraints would make it difficult to meet the full shortfall in natural gas imports from LNG alone.[50]

As with the Fukushima counter-factual, I consider a number of alternative contracting regimes for the LNG industry. The baseline regime involves no restrictions on long-term contracting. The "No Contracts" regime assumes that all of trade is carried out via the spot market. "No European Contracts" assumes that only the long-term LNG contracts agreed to by European buyers are non-binding, but all other contracts are binding. "No non-European Contracts" assumes instead that

[50]The binding constraint is not so much spare regasification capacity, which is plentiful in European countries, but rather constraints in inter-connection capacity. This makess it challenging to deliver natural gas from European countries with LNG import capabilities (such as Spain, France or the UK) to other European countries reliant on Russian natural gas who do not have their own LNG import facilities (such as Germany). See IEA (2022) for further discussion of this point.

only the long-term contracts agreed to by European buyers are binding, but all other contracts are binding. Finally, I compare all of these regimes with a "Fully Competitive" regime (equivalently, a regime where no destination clauses are used), where all LNG is competitively allocated, which provides the efficient benchmark against which we can compare all the other contracting regimes.

Table C15: Effect of of a hypothetical shutdown of Russian natural gas exports to Europe on LNG trade and industry surplus

	Baseline	No Contracts	No Europ. Contracts	No non-Europ. Contracts	Fully Competitive
Welfare ($bn)	205.7	207.4	205.5	207.9	211.8
Seller Surplus ($bn)	96.0	98.1	99.9	95.2	96.1
Buyer Surplus ($bn)	109.8	109.3	105.6	112.7	115.7
Europe's total LNG imports (mt)	86.9	89.8	84.8	92.1	100.0
Increase due to shutdown (mt)	46.8	49.8	44.7	52.1	60.0

Note: Welfare, seller surplus and buyer surplus are reported for the year 2017. The fourth row reports Europe's total LNG imports in 2017. The fifth row reports the increase in Europe's total LNG imports induced by the shutdown of Russian natural gas exports to Europe, by subtracting from the fourth row Europe's total LNG imports in the baseline scenario with no shutdown. I assume that the shutdown of Russian natural gas exports cause demand for natural gas to increase by 100 bcm.

Table C14 illustrates the results. The industry responds more efficiently to the increase in European LNG demand when the industry does not use long-term contracts: welfare is $205.7bn in the baseline regime, but rises to $207.4 bn in the "No Contracting" regime, a $1.65bn welfare gain. Sellers end up selling more LNG to Europe if they are not bound by long-term contracts, with European LNG imports increasing by 49.8 mt in the "No Contracts" scenario versus 46.8 mt in the baseline regime (a difference of 3 mt). Just as we saw with the analysis of the Fukushima demand shock, the welfare gains come from eliminating contracts of non-European buyers who are not subject to the demand shock, thus freeing up LNG supplies that can be sold to European buyers instead. Eliminating long-term contracts of non-European buyers leads to welfare gains and more LNG allocated to Europe, whereas removing European contracts actually leads to a smaller welfare loss, due to increased market power on the spot market. Finally, the competitive allocation would lead to a much bigger increase in European LNG demand (60 mt, as opposed to 46.8 mt in the baseline) and a $6.1 bn increase in welfare.

Additional References

BACHMANN, R., D. BAQAEE, C. BAYER, M. KUHN, A. LÖSCHEL, B. MOLL, A. PEICHL, K. PITTEL, AND M. SCHULARICK (2022): "What if? The economic effects for Germany of a stop of energy imports from Russia," *ECONtribute Policy Brief*, 28, 2022.

HENDEL, I. AND A. NEVO (2006): "Measuring the Implications of Sales and Consumer Inventory Behavior," *Econometrica*, 74, 1637–73.

IEA (2022): "A 10-Point Plan to Reduce the European Union's Reliance on Russian Natural Gas," .

KAUFMANN, D., A. KRAAY, AND M. MASTRUZZI (2004): "Governance matters III: Governance indicators for 1996, 1998, 2000, and 2002," *The World Bank Economic Review*, 18, 253–287.

MILLER, N. AND M. OSBORNE (2014): "Spatial differentiation and price discrimination in the cement industry: evidence from a structural model," *RAND Journal of Economics*, 45, 221–247.

MIYAMOTO, A., M. NAKAMURA, AND C. ISHIGURO (2012): "A realistic perspective on Japan's LNG demand after Fukushima," Working paper, Oxford Institute for Energy Studies.

NEIDELL, M. J., S. UCHIDA, AND M. VERONESI (2019): "Be Cautious with the Precautionary Principle: Evidence from Fukushima Daiichi Nuclear Accident," Working paper, National Bureau of Economic Research.

NEVO, A. AND F. ROSSI (2008): "An Approach for Extending Dynamic Models to Settings with Multi-Product Firms," *Economics Letters*, 100, 49–52.

NUNN, N. (2007): "Relationship-specificity, incomplete contracts, and the pattern of trade," *The Quarterly Journal of Economics*, 122, 569–600.

PESCATORI, M. A., M. J. C. BLUEDORN, M. C. KOCH, S. ALBRIZIO, AND M. STUERMER (2022): "Market Size and Supply Disruptions: Sharing the Pain of a Potential Russian Gas Shutoff to the European Union," Working paper, International Monetary Fund.

RYAN, S. (2012): "The Costs of Environmental Regulation in a Concentrated Industry," *Econometrica*, 80, 1019–1061.

STAPCZYNSKI, S. AND D. MURTAGH (2019): "The Future Is Now for LNG as Derivatives Trading Takes Off," *Bloomberg*.

TERAZONO, E. (2019): "Looser restrictions prompt surge in LNG liquidity," *Financial Times*.

Supplementary Material for "Long-term contracts and efficiency in the liquefied natural gas industry"

Nahim Bin Zahur

December 2022

D Model Simulations

In this section, I describe two sets of Monte Carlo simulations of the model developed in Section 4. The simulations use simplified versions of the model and are designed to highlight specific mechanisms of the model. Firstly, I explore how bargaining power and outside options affects the contracting and investment decisions of the parties (Section D.1). Second, I explore the short-run allocation effects of using contracts, as well as the size and nature of contracting externalities (Section D.2.1).

D.1 Bargaining power, investment and contracting

As discussed in Section 4.4, the relative bargaining power and the presence of outside options affect the return on investment earned by the seller. This in turn affects the incentives of the seller and buyer to sign contracts ex-ante, in Stage 1. Here I further elaborate on these properties and describe Monte Carlo simulations that illustrate the key features of the model.

Bargaining power, outside options and (under-)investment: I begin with how bargaining power influences the seller's investment decision. For simplicity, consider the case where a single seller bargains with a single buyer in Stage 3. The level of investment that maximizes the joint surplus of the seller and buyer (which I will call the efficient bilateral investment level) is:

$$K_i^{**} = argmax_{K_i}\big[\underbrace{V_i^3(q_i^{c,1},q_i^{c,3},K_i,Y_{-i})+W_j^3(q_i^{c,1},q_i^{c,3},K_i,Y_{-i},\eta_{ij}^3)}_{\text{Sum of seller and buyer surplus}}-\Gamma_i(K_i,\eta_i^2)\big] \tag{D1}$$

By contrast, the investment actually chosen by the seller maximizes the sum of the seller's payoff from investing and the lump-sum transfer received from the buyer. Combining equations (6) and (7) yields the following equation for the seller's choice of investment:

$$\begin{aligned} K_i^{*} = argmax_{K_i}\big[\tau\big(&\underbrace{V_i^3(q_i^{c,1},q_i^{c,3},K_i,Y_{-i})+W_j^3(q_i^{c,1},q_i^{c,3},K_i,Y_{-i},\eta_{ij}^3)}_{\text{Sum of seller and buyer surplus}}\big)-\Gamma_i(K_i,\eta_i^2)+ \\ &+(1-\tau)\underbrace{V_i^3(q_i^{c,1},q_{i,\setminus ij}^{c,3},K_i,Y_{-i})}_{\text{Seller's disagreement payoff}}-\tau\underbrace{W_j^3(q_i^{c,1},q_{i,\setminus ij}^{c,3},K_i,Y_{-i})}_{\text{Buyer's disagreement payoff}}\big] \end{aligned} \tag{D2}$$

Comparing equations (D1) and (D2) illustrates how the seller's actual investment K_i^* differs from the efficient bilateral investment level K_i^{**}. Consider first the special case where the investment has no value to either the seller or buyer outside of their contractual relationship, so that their outside options (i.e., disagreement payoffs) do not depend on K_i. This is the hold-up effect: the seller bears the full cost of investment, but only enjoys a share τ of the return on investment, and as such under-invests.[51] The smaller the seller's Nash bargaining weight τ, the more severe the under-investment.

In practice, though, investment is valuable to the seller and buyer even outside their contractual relationship, since in the event of disagreement, they are still able to participate in the spot market, where they both benefit from the added investment by the seller. Thus the extent of under-investment will also depend on the relative strength of the outside options of the seller and buyer, with the two having opposite effects on investment. The seller's outside option $V_i^3(q_i^{c,1}, q_{i,\setminus ij}^{c,3}, K_i, Y_{-i})$ is increasing in K_i (since by investing more they can increase their spot market profits). Thus, the presence of an outside option for the seller ameliorates the problem of under-investment, and leads to greater investment by the seller. On the flip side, though, the buyer's outside option $W_j^3(q_i^{c,1}, q_{i,\setminus ij}^{c,3}, K_i, Y_{-i})$ also increases in K_i, since additional capacity on the spot market lowers the expected price paid by the buyer for spot purchases. Thus, the presence of an outside option for the buyer worsens the problem of under-investment, and leads to lower investment.

The presence of these outside options does not entirely eliminate under-investment, however. K_i^* will only equal the efficient bilateral investment level K_i^{**} if the seller has complete bargaining power and the ability to fully capture the surplus from trade when negotiating with the buyer. This requires not only that the Nash bargaining weight τ equals 1 (meaning the seller has the ability to make take-it-or-leave-it offers), but also that the buyer's disagreement payoff $W_j^3(q_i^{c,1}, q_{i,\setminus ij}^{c,3}, K_i, Y_{-i})$ is non-increasing in K_i (meaning that the seller is able to deny the buyer any benefits from their investment in the event of disagreement). By contrast, provided the buyer has some bargaining power (e.g., if $\tau < 1$ or if the buyer has a non-trivial outside option that is rising in K_i), the seller will under-invest, since the seller does not fully internalize the benefit that the buyer enjoys from the investment.

Because of the potential for under-investment, the seller and buyer have an incentive to sign a larger contract ex-ante (in Stage 1) to forestall under-investment, as discussed in Section 4.4.

Simulations of investment and contracting model In the remainder of this section, I report the results from Monte Carlo simulations to explore the relationship between bargaining power, investment and contracting. The simulations are based on a simplified version of the model where a single seller interacts with a single buyer, and both the seller and the buyer have access to a spot

[51] Mathematically, under-investment arises since the actual investment level K_i^* only maximizes $\tau(V_i^3 + W_j^3) - \Gamma_i$, whereas the optimal investment level K_i^{**} would maximize $(V_i^3 + W_j^3) - \Gamma_i$.

market (that forms their "outside option" in case the bargaining breaks down). The seller and buyer play the three-stage game outlined in Section 4: they can sign a contract prior to investment (Stage 1); the seller then chooses how much to invest (Stage 2); after the seller has already committed to the investment, the seller and buyer can get together again to sign a new contract, on top of the existing contract they already signed (Stage 3). Finally once contracts are signed and investments made, the seller sells any excess capacity on the spot market, and the buyer purchases any excess LNG requirement on the spot market.

Consistent with the theoretical predictions, we find that if the seller and the buyer cannot contract in Stage 1, and as long as the buyer has some bargaining power, the seller will tend to under-invest relative to the optimum. The larger the Nash bargaining weight of the seller, the less severe the under-investment and the higher the level of capacity the seller builds, as illustrated in Figure D1. Similarly, the weaker the seller's outside option (as captured by an increasing cost of selling to the spot market), the more severe the under-investment and the smaller the level of capacity the seller builds (Figure D2).[52]

Next, if we allow the seller and buyer to contract before investment, we find that the ability to sign ex-ante contracts prior to investment increases investment and raises welfare. Moreover, the lower the Nash bargaining weight of the seller, the more contracting takes place in Stage 1 and the less contracting takes place in Stage 3 (see Figure D3). Similarly, the weaker the outside option of the seller, the more contracting takes place in Stage 1 and the less contracting takes place in Stage 3 (see Figure D4). This is because the under-investment in Stage 2 is more severe when sellers have low bargaining power relative to the buyers, and anticipating this the seller and buyer sign larger contracts in Stage 1.

D.2 Long-term contracts, allocative efficiency and contracting externalities

The next set of Monte Carlo simulations focuses on the short-run allocation effects of using long-term contracts, as well as externalities imposed by the use of long-term contracts. For this analysis, I rely on the model of the spot market developed in Section 4.2. The main purpose of the simulations is to look at how contracts affect equilibrium allocations and prices in the Cournot model, and identify the conditions under which long-term contracts could increase/reduce allocative efficiency (Section D.2.1). In Section D.2.2, I then explore the conditions under which contracting externalities can emerge. For tractability, I do not endogenize the choice of contracts, focusing instead on how different contract quantities affect the equilibrium allocations, prices and welfare.

Throughout the simulations I assume there are 2 sellers and 2 buyers. Seller 1 is located closer to Buyer 1 than Seller 2 ($d_{11} < d_{12}$), and Seller 2 is located closer to Buyer 2 than Seller 1 ($d_{22} < d_{21}$).

[52] I also find, in a similar vein, that the stronger the buyer's outside option, the more severe the under-investment. The detailed results are available upon request.

D.1.1 Bargaining power and outside options: effect on (under-)investment

Figure D1: Effect of seller's Nash bargaining weight τ on K

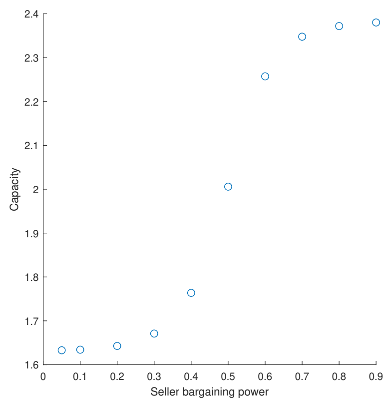

Figure D2: Effect of seller's cost of selling to spot market on K

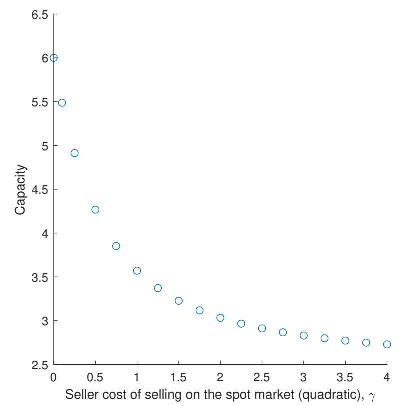

D.1.2 Bargaining power and outside options: effect on ex-ante contracting

Figure D3: Share of contracts signed in Stage 1 (ex-ante) vs. seller bargaining power

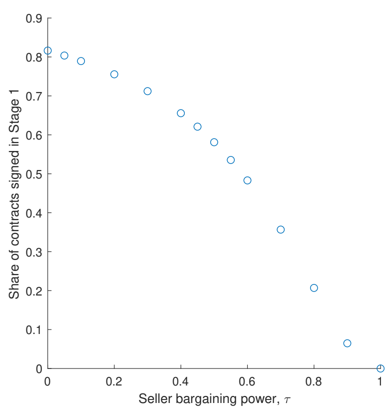

Figure D4: Share of contracts signed in Stage 1 (ex-ante) vs. seller's cost of selling to spot market

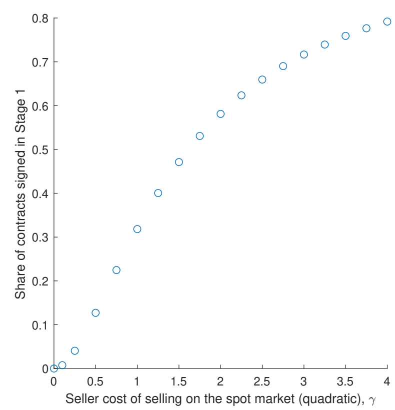

I assume that firms cannot produce beyond their capacity K_i and have constant marginal costs for any output below K_i. I set the marginal cost of production to equal 4. I assume that demand shocks in the two buying countries are uniformly and independently distributed: $\varepsilon_j \sim U[D_l, D_h]$,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

I consider two different simulations (see Table D2):

1. S1, where firms have lots of capacity relative to demand and always produce below capacity.

2. S2, where firms have very limited capacity will always produce at full capacity (even when they have no contracts).

Table D1: Assumptions maintained in all scenarios

Demand Slope	$b = 1$	
Demand Shock	$\varepsilon_j \sim U[d_l, d_h]$	
Demand Parameters	$D_l = 40$	$D_h = 80$
Marginal cost of production	4	
Shipping costs from seller 1	$d_{11} = 1$	$d_{12} = 5$
Shipping costs from seller 2	$d_{21} = 5$	$d_{22} = 1$

Table D2: Simulations

	S1	S2
Description:	High capacity	Limited capacity
K_1	100	15
K_2	100	15

D.2.1 Allocative efficiency

To investigate how contracts affect allocations, I consider a variety of contract configurations (Table D3). I first look at the case where firms sign no contracts (C1). In cases C2-C5, I then progressively increase the contract quantity signed by each seller, making sure that each seller has the same total contracted quantity. In each case, I assume that each seller signs 62.5% of their total contract quantity with the nearest buyer, and 37.5% with the faraway buyer; the reason is that in the scenario with no contracts, I find that sellers on average sell 62.5% of their output to the nearest buyer.

Scenario 1 (S1): firms always produce below capacity: Table D4 shows how the allocation changes as we introduce long-term contracts. When firms are unconstrained by capacity, the higher the contracted level of quantity, the higher is total production, the lower are spot prices and the higher is welfare. This confirms that contracts are welfare-improving when firms are unconstrained by capacity, just as in Allaz and Vila (1993). The reason is that the greater the contracted quantity, the more competitively the firm behaves on the spot market since they can do so without reducing

Table D3: Contract configurations

		Seller 1 to Buyer 1 q^c_{11}	Seller 1 to Buyer 2 q^c_{12}	Seller 2 to Buyer 1 q^c_{21}	Seller 2 to Buyer 2 q^c_{22}	Seller 1 Total $q^c_{11} + q^c_{12}$	Seller 2 Total $q^c_{21} + q^c_{22}$
No contracts	C1	0	0	0	0	0	0
Sellers sign contracts	C2	1.25	0.75	0.75	1.25	2	2
	C3	4.37	2.63	2.63	4.37	7	7
	C4	8.75	5.25	5.25	8.75	14	14
	C5	9.375	5.625	5.625	9.375	15	15

the payoff they receive on their contracted output.

Table D4: Effect of contracts when firms produce below capacity

			Prices		Trade flows					
					Seller 1		Seller 2		Total	Total
	Contracts	Welfare	p_1	p_2	q_{11}	q_{12}	q_{21}	q_{22}	Prod.	Capacity
Competitive	0	3152	5	5	54.6	0	0	55.4	110	200
Cournot, C1	0	2626	24.5	24.8	19.5	15.8	15.5	19.8	71	200
Cournot, C2	4	2651	23.9	24.1	20.1	15.9	15.6	20.4	72	200
Cournot, C3	14	2710	22.2	22.5	21.6	16.1	15.8	21.8	75	200
Cournot, C4	28	2783	19.9	20.1	23.6	16.4	16.1	23.9	80	200
Cournot, C5	30	2792	19.5	19.8	23.9	16.4	16.1	24.2	81	200

For each scenario, I draw 1000 different pairs of demand shocks $(\varepsilon_1, \varepsilon_2)$*. I solve for spot prices and allocation, taking any contracted flows as given. 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.*

Scenario 2 (S2): firms always produce at capacity: Table D5 looks at the effect of contracts when firms have limited capacity. Now contracts no longer affect total production, so the pro-competitive effect of contracts identified by Allaz and Vila (1993) is less strong. Moreover when the contracted quantities are sufficiently large in relation to the available capacity, welfare decreases as we increase contracts. This shows that contracts are welfare-reducing when firms are fully capacity-constrained.

Table D5: Effect of contracts when firms produce at capacity

			Prices		Trade flows					
					Seller 1		Seller 2		Total	Total
	Contracts	Welfare	p_1	p_2	q_{11}	q_{12}	q_{21}	q_{22}	Prod.	Capacity
Competitive	0	1466	45	45	12.5	3	2	12.8	30	30
Cournot, C1	0	1438	44.8	45.1	9.4	5.6	5.4	9.6	30	30
Cournot, C2	4	1439	44.8	45.1	9.6	5.4	5.2	9.8	30	30
Cournot, C3	14	1436	44.8	45.2	9.7	5.3	5.0	10.0	30	30
Cournot, C4	28	1391	44.6	45.4	9.3	5.7	5.6	9.4	30	30
Cournot, C5	30	1379	44.6	45.4	9.4	5.6	5.6	9.4	30	30

For each scenario, I draw 1000 different pairs of demand shocks $(\varepsilon_1, \varepsilon_2)$*. I solve for spot prices and allocation, taking any contracted flows as given. 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.*

What is the source of these welfare losses? In order to understand this better, I look separately at scenarios with symmetric and asymmetric demand shocks in Table D6 and Table D7. As we might expect, the welfare losses from contracts happen exactly when demand is high in one market but low in the other market. The spot market is able to respond to the demand asymmetry by re-directing LNG to the market with higher demand, but contracts are unable to do this. Table D6 shows that the use of contracts lead to a marked reduction in welfare when demand is high in market 1 and low in market 2. This is because sellers sell contracted LNG to buyer 2 even though buyer 1 is willing to pay a higher spot price than buyer 2. This also means that price differentials get larger because of contracts: comparing the second row with the last row, we can see that reducing contracted quantities would lead to smaller price differences across regions. (The results are similar when demand is low in market 1 but high in market 2). By contrast, as Table D7 confirms, contracts have a much more modest effect on welfare when demand is high in both markets (and this is also true when demand is low in both markets).

Table D6: Effect of contracts when demand is high in market 1, low in market 2

			Prices		Trade flows					
					Seller 1		Seller 2			
	Contracts	Welfare	p_1	p_2	q_{11}	q_{12}	q_{21}	q_{22}	Total prod.	Total capacity
Competitive	0	1579	47	43	15.0	0	12	2.8	30	30
Cournot, C1	0	1565	49.7	40.1	14.2	0.8	10.3	4.7	30	30
Cournot, C2	4	1562	49.9	39.9	14.1	0.9	10.3	4.7	30	30
Cournot, C3	14	1533	51.6	38.2	12.4	2.6	10.3	4.7	30	30
Cournot, C4	28	1403	58.3	31.5	9.8	5.3	6.3	8.8	30	30
Cournot, C5	30	1377	59.3	30.5	9.4	5.6	5.6	9.4	30	30

For each scenario, I draw 1000 different pairs of demand shocks ($\varepsilon_1, \varepsilon_2$). I solve for spot prices and allocation, taking any contracted flows as given. The table presents averages across realizations of ε where $\varepsilon_1 > 70$ and $\varepsilon_2 < 50$.

Table D7: Effect of contracts when demand is high in both market 1 and market 2

			Prices		Trade flows					
					Seller 1		Seller 2			
	Contracts	Welfare	p_1	p_2	q_{11}	q_{12}	q_{21}	q_{22}	Total prod.	Total capacity
Competitive	0	1862	59	60	14.8	0	0	14.9	30	30
Cournot, C1	0	1821	59.4	59.7	9.4	5.6	5.4	9.6	30	30
Cournot, C2	4	1823	59.4	59.7	9.6	5.4	5.1	9.9	30	30
Cournot, C3	14	1828	59.4	59.7	10.3	4.7	4.5	10.5	30	30
Cournot, C4	28	1820	59.3	59.8	9.6	5.4	5.3	9.7	30	30
Cournot, C5	30	1817	59.2	59.9	9.4	5.6	5.6	9.4	30	30

For each scenario, I draw 1000 different pairs of demand shocks ($\varepsilon_1, \varepsilon_2$). I solve for spot prices and allocation, taking any contracted flows as given. The table presents averages across realizations of ε where $\varepsilon_1 > 70$ and $\varepsilon_2 > 70$.

D.2.2 Contracting externalities

I next investigate whether contracts impose externalities on other firms. I take as the baseline case the Courot game with no contracts (the "Cournot, C1"). I then look at what happens to the welfare of all four firms when (i) seller 1 signs a contract with buyer 1 (their nearby buyer) (ii) seller 1 signs a contract with buyer 2 (their farway buyer). I do so both in the scenario where sellers have excess capacity (Scenario 1), and when firms have limited capacity (Scenario 2).

Scenario 1 (S1): firms always produce below capacity: First, consider the effect of a contract signed between seller 1 and buyer 1 when sellers are not capacity-constrained (Table D8). Because of market power on the spot contract, this contract is beneficial to both seller 1 and buyer 1. Seller 1 benefits from the contract because by committing in advance to a long-term contract, it is able to increase its market share: this is the Stackelberg effect of signing quantity contracts that was identified by Allaz and Vila (1993). Buyer 1 benefits because in equilibrium, it receives a higher quantity at a lower price. This shows that when sellers exercise market power on the spot market, sellers and buyers have incentives to contract ex-ante: intuitively, the deadweight loss from market power means there are "gains on the table" from increasing the quantity sold to the buyer, which the parties can exploit by signing a forward contract.

How does this contract affect the two excluded parties: seller 2, and buyer 2? Because the sellers are not capacity-constrained, then their quantity decisions across markets are independent. As such, buyer 2 is unaffected by the contract signed by seller 1 and buyer 1, so the externality imposed on buyer 2 is zero. However, seller 2 now faces stronger competition when selling to buyer 1, so the contract makes seller 2 worse off. Similarly, a contract signed between seller 1 and buyer 2 makes seller 2 worse off, without affecting buyer 1. In the absence of capacity constraints, therefore, contracts impose a negative externality on sellers who do not sign the contracts but compete directly with the sellers who do, but do not impose any externalities (positive or negative) on other buyers.

Table D8: Effect of contract signed by seller 1, when firms produce below full capacity

	Contract quantity				Change in welfare, relative to baseline				
	q^c_{11}	q^c_{12}	q^c_{21}	q^c_{22}	Seller 1	Seller 2	Buyer 1	Buyer 2	Externality
Nearby contracting	1	0	0	0	6.28	-10.23	11.73	0.00	-10.23
Faraway contracting	0	1	0	0	5.04	-13.09	0.00	11.92	-13.09

Scenario 2 (S2): firms always produce at capacity: Next, we repeat this analysis when firms are capacity constrained. This leads to an entirely different set of results, as shown by Table D9. The key difference now is that because sellers are capacity constrained, their decisions in different

markets are inter-related. If they raise their sales to one buyer, that must be accompanied by reduced sales to other buyers.

As a consequence, buyers are now in direct competition with each other. A buyer benefits from signing a contract, but at the expense of the other buyer. the first row of Table D9 shows that if buyer 1 contracts with seller 1, buyer 1 is better off while buyer 2 is worse off. Conversely, if buyer 2 contracts with seller 1, buyer 2 is better off while buyer 1 is better off.

Intuitively, this is because by signing a contract, a buyer is able to (weakly) increase the total quantity they are able to purchase, since the contract provides a floor below which their purchases will not drop. Contracts also protect buyers against sellers' exercise of market power, since the seller with whom the buyer contracts now has an incentive to compete more aggressively when selling to that buyer (due to the Allaz and Vila (1993) effect). As the second panel of Table D9 shows, when a buyer signs a contract, the average quantity that they purchase increases, and the average price they pay decreases.

Table D9: Effect of contract signed by seller 1, when firms produce at capacity

	Contract quantity				Change in welfare, relative to baseline				
	q^c_{11}	q^c_{12}	q^c_{21}	q^c_{22}	Seller 1	Seller 2	Buyer 1	Buyer 2	Externality
Nearby contracting	1	0	0	0	0.50	1.49	2.42	-2.56	-1.07
Faraway contracting	0	1	0	0	-0.86	-1.26	-2.59	2.61	-3.85

	Change in quantity traded				Change in prices		Change in buyer total Q	
	q_{11}	q_{12}	q_{21}	q_{22}	Buyer 1	Buyer 2	Buyer 1	Buyer 2
Nearby contracting	0.33	-0.33	-0.16	0.16	-0.17	0.17	0.17	-0.17
Faraway contracting	-0.34	0.34	0.17	-0.17	0.17	-0.17	-0.17	0.17

But through the exact opposite of these forces, buyers are negatively impacted by contracts signed by other buyers, since it reduces the quantity they can purchase (especially in states of the world where they have high demand), and may worsen the market power dynamics, since contracts signed by other sellers may increase the monopoly power of other, uncontracted sellers. As the second panel of Table D9 shows, when a buyer signs a contract, the average quantity purchased by the *rival* buyer decreases, and the average price increases. The intuition behind these findings is similar to Bolton and Whinston (1993), who find that when a seller is capacity-constrained and can sell to multiple buyers (but is unable to meet the demand of every buyer), vertical integration between the seller and one of the buyers benefits that buyer at the expense of the other buyers.

The effects of contracts on welfare of the sellers is more complicated: unlike the buyers who are price-takers on the spot market, sellers are strategic actors in the spot market, and contracts affect their spot market decisions. As the first row of the first panel of Table D9 illustrates, if a seller (in this case seller 1) signs a contract with their nearby buyer, then *both* sellers benefit. In this case, the contract imposes a positive externality on the excluded seller. But if the seller signs a contract with their faraway buyer (second row of the first panel of Table D9), then both sellers are made worse off: in this case, the contract imposes a negative externality on the other seller.

These results, while counter-intuitive at first glance, stem from the nature of the market power in spatially differentiated markets with capacity constraints. With binding capacity constraints, market power does not affect the total production of each firm, but instead distorts their allocation of output across markets, with each firm selling too large a share of their output to their faraway buyer. This outcome is undesirable for both sellers, since they end up incurring higher shipping costs. Both sellers would be better off if they could jointly agree to reduce their sales to the faraway buyer and instead increase their sales to the their nearby buyer, but no one firm has an incentive to unilaterally deviate from the Cournot equilibrium. In other words, the Cournot equilibrium involves firms competing too strongly in the markets of their rivals, which hurts both firms.

Quantity contracts increase the competition faced by sellers in the market where the buyer has signed a contract, since the seller on that contract behaves more competitively. For example if seller 1 and buyer 1 sign a contract, seller 1 will compete more aggressively in market 1, so seller 2 can expect to receive a lower price in market 1. Conversely, quantity contracts decrease the competition faced by sellers in the market where buyers have not signed a contract, by reducing the total quantity available in that market. Going back to the example where seller 1 and buyer 1 sign a contract, since seller 1 competes more aggressively in market 1, they must compete less aggressively in market 2 (due to capacity constraints), meaning seller 2 can now expect to receive a higher spot price in market 2.

Consequently, the effects of contracts on seller welfare depend on how they influence the level of competition faced by sellers in nearby and distant markets. When seller 1 signs a contract with their nearby buyer 1, then seller 1 is induced to increase their sales to buyer 1, which induces seller 2 to increase their sales to buyer 2. In other words, contracts signed with nearby buyers induce sellers to intensify competition in their nearby markets and reduce the extent to which they compete in faraway markets, which ameliorates the Prisoner's Dilemma and benefits both sellers. By contrast, when seller 1 signs a contract with their faraway buyer 2, seller 1 is induced to raise their sales to buyer 2, which now induces seller 2 to increase their sales to buyer 1 (seller 1's nearby buyer). In other words, contracts signed with faraway buyers induce sellers to intensify competition in the faraway markets, which worsens the Prisoner's Dilemma and makes both sellers both worse off.

While seller externalities are ambiguous in sign, buyer externalities are unambiguously negative, and this implies that the net externality from contracting is negative. In the last column of the top panel of Table D9, we can see that contracts *on net* reduce the welfare of agents not party to the contract, regardless of whether the seller contracts with their nearby or faraway buyer.

Summary: These simulations suggest that in spatially differentiated markets where sellers exercise market power and may be capacity constrained, contracts impose externalities on both excluded buyers and excluded sellers who are not party to the contract. Contracts impose negative externalities on excluded buyers if sellers are capacity-constrained, since they reduce the quantity available for these buyers to purchase. Contracts may impose negative or positive externalities on excluded sellers: the externalities are negative if the contracts lead the seller to face greater competition in their major market, but the externalities are positive if they reduce competition faced by the seller in their main market. Regardless of whether or not seller externalities are positive, I find that the net external effects of contracts are negative: that is, the total welfare of excluded parties decreases when a contract is signed.

E Further Estimation Details

E.1 Approximations of seller and buyer expected payoffs

I employ parametric approximations of seller and buyer expected payoffs (as discussed in Section 5.3 in page 30). Here I provide more details on these approximations.

I assume that each seller's payoffs can be approximated by a set of L_s basis functions $u_1, ..., u_{L_s}$, and each buyer's expected payoffs can be approximated by a set of L_b basis functions $\phi_1, ..., \phi_{L_b}$:

$$\pi^s_{it}(q^c_t, K_t) \simeq \sum_{l=1}^{L_s} b^s_l u_l(q^c_t, K_t, x_t) \tag{E3}$$

$$\pi^b_{jt}(q^c_t, K_t) \simeq \sum_{l=1}^{L_b} b^b_l \phi_l(q^c_t, K_t, x_t) \tag{E4}$$

where b^s_l and b^b_l are unknown approximating parameters that need to be estimated.

In order to estimate the approximating parameters, I first randomly draw S combinations of q^c_t (vector of contracted quantities) and K_t (capacities). For each of these S draws of the states, I randomly draw D realizations of ε_t (demand shocks). For each state and demand draw, I then solve for the spot market equilibrium in order to obtain per-period payoffs to buyers and sellers. I then take the expectation over demand shocks in order to get per-period *expected* payoffs.

I am then left with S simulations of the spot market, where for each simulation I know q^c_t, K_t, $\pi^s_{it}(q^c_t, K_t, x_t)$ for each seller i and $\pi^b_{jt}(q^c_t, K_t, x_t)$ for each buyer j. I now regress $\pi^s_{it}(q^c_t, K_t, x_t)$ on basis functions of (q^c_t, K_t, x_t) in order to obtain b^s_l and regress $\pi^b_{jt}(q^c_t, K_t, x_t)$ on basis functions of (q^c_t, K_t, x_t) in order to obtain b^b_l.

Implementation: When implementing the above procedure, I carry out $S = 6000$ simulations of the spot market, each with a different draw of capacity and contracts, and set $D = 200$ (meaning 200 different draws of the demand shocks for each of the 6000 simulations).[53] This requires solving the spot market equilibrium 1.2 million times, which is computationally intensive. However, this only needs to be done once prior to estimation, and parallel computation can be used to speed up the process. After integrating out the demand shocks, I am left with 6,000 sets of simulations of the spot market, with expected payoffs for every buyer and every seller. There are multiple buyers and sellers, so I have 156,000 observations of buyer expected payoffs to estimate buyer approximation parameters, and 79,987 observations of seller expected payoffs.

The choice of basis functions is crucial in ensuring well-behaved approximations. In order to keep the problem computationally tractable, I make a few simplifying assumptions. I first assume that firms do not keep track of the state variables of each one of their rivals. Instead, firms only keep track of two sufficient statistics representing the choices of their rivals: the total capacity of rivals, and the total contract quantity signed by all rival firms. As discussed in Section 5.3, this approach has precedents in the literature on dynamic games estimation, most notably by Weintraub et al. (2008) and Benkard et al. (2015) who propose the notion of oblivious equilibrium as a way to approximate Markov perfect equilibrium in dynamic models.[54]

With the above assumption, I am able to reduce the dimension of the *endogenous* state variables (i.e., capacity K_t and contract quantities q_t). However, the dimension of x_t (which includes the *exogenous* variables) is very large, since it includes demand shifters for every buyer in the market, and the distance matrix summarizing nautical distances between every seller and buyer in the market. Including polynomial functions of each component of x_{it} would lead to a very large number of basis functions, which can result in poorly behaved approximations.

As such, I make a second simplifying assumption that is very much in the same spirit of the first assumption: firms do not keep track of every element of x_t, and instead only keep track of selected indices that succinctly summarize the effect of x_t on their payoffs. This approach of collapsing high-dimensional firm-level state variables into low-dimensional indices is similar to Hendel and Nevo (2006) and Nevo and Rossi (2008), who use the "inclusive value" to capture the impact of changing product attributes on future profits.

To motivate my choice of indices to summarize the effect of x_t on payoffs, it is useful first to revisit the demand curve. The demand for buyer j in period t (equation (10)) can be rewritten as:

$$\begin{aligned} Q_{jt} &= -bR_{jt}p_{jt} + \alpha R_{jt} + \theta_j R_{jt} + x_{jt}R_{jt}\theta_{dx} + \varepsilon_{jt}R_{jt} \\ &= -bR_{jt}p_{jt} + \underbrace{z_{jt}}_{\text{Persistent component of demand}} + \underbrace{\hat{\varepsilon}_{jt}}_{\text{Demand shock}} \end{aligned}$$

[53] I have data on the spot market for 12 years between 2006 - 2007, each with a different set of sellers and buyers. The 6,000 simulations were evenly divided across these 12 years, to ensure that the simulations were representative of the observed sample.

[54] Empirical applications of this approach include Sweeting (2015), Gerarden (2017), Jeon (2018) and Chen and Xu (2021).

where $z_{jt} = \alpha R_{jt} + \theta_j R_{jt} + x_{jt} R_{jt} \theta_{dx}$ can be thought of as the persistent component of demand, which is a function of the buyer's size R_{jt}. buyer fixed effects θ_j and buyer characteristics x_{jt}.

As z_{jt} increases, the buyer's demand curve for LNG shifts outwards, which will change buyers' consumer surplus. Thus, I include z_{jt} when approximating buyer payoffs, as a way to summarize how exogenous state variables shift buyer j's demand. Likewise, when approximating seller payoffs, I include the variable $z_j = \sum_t z_{jt}$, which is a measure of aggregate demand for LNG.

In addition to these demand indices, I also construct a separate set of indices that capture the effect of geography on seller and buyer payoffs. Buyers who are located far away from sellers will be forced to pay higher spot prices on average (to compensate sellers for the higher shipping cost they have to incur), so their expected consumer surplus will be on average lower. I capture this effect by allowing buyer's payoffs on the average distance from buyer j to sellers ($\bar{d}_j$). Sellers who are located far from most buyers (especially buyers with high demand) will similarly have a lower expected payoff. To capture this effect, I include $\bar{d}^z_{it}$ when approximating seller's payoff, which is the seller's average distance from other buyers weighted by the demand state z_{jt} of each buyer.

Buyer payoff approximation: Putting together all these assumptions, buyer expected payoffs depend on their total contracted quantity $Q^c(j)$, rival contracted quantity ("rival Q^c"), total capacity in that year K, as well as the two indices that summarize the role of exogenous states (demand state z_{jt} and average distance $\bar{d}_j$). I model buyers' expected payoffs as quadratic functions of $Q^c(j)$, "rival Q^c", and K, with interactions of those terms with z_{jt} and $\bar{d}_j$. The approximation parameters can be estimated by an OLS regression of buyer expected payoffs on these selected basis functions, using the simulated sample.[55]

Seller payoff approximation: Sellers' expected payoffs are functions of their total contracted quantity Q^c_i, their capacity K_i, rival contracted quantity ("rival Q^c) and rival capacity in that year ("rival K"), as well as aggregate demand z_t and the seller's demand-weighted average distance from other buyers $\bar{d}^z_{it}$. Like with buyers, I assume sellers' payoffs include a quadratic function of the endogenous variables plus interaction terms between each of the endogenous variables and the two indices summarizing the exogenous variables.

But the quadratic basis functions do not capture capacity constraints all that well: as we know, sellers' marginal costs rise steeply in their capacity utilization. As such, I also include a basis function especially designed to capture capacity constraints, $log(\psi K_i - Q^c_i)$ (where ψ is a tuning parameter). The idea behind this is that as Q^c_i gets larger (relative to K_i), the seller's capacity utilization is higher and the seller is more likely to be highly capacity-constrained, which will lower their expected payoffs since the marginal costs of operating are higher when capacity utilization is

[55]These results are available upon request.

high. This idea is easiest to see when we set $\psi = 1$: in that case, the closer the seller's contracted quantity gets to their capacity, the smaller the value of the basis function (as the seller's marginal cost rises). In the extreme, as Q_i^c approaches K_i, the basis function approaches negative infinity, which would mean sellers' payoffs would also approach negative infinity as they experience the high marginal costs of operating at full capacity. Since in practice I do observe sellers occasionally signing contracts amounting to close to 100% of their capacity, I set ψ to be slightly larger than 1, to permit this behavior.[56]

E.2 Comparison of model-predicted contract prices with contract prices inferred from customs data

LNG contract pricing formulas are confidential, and systematic data on pricing formulas is not available. This is why this paper takes the approach of estimating the structural parameters without using any information at all on contract prices. However, in reality (despite contract confidentiality) some information *is* available on LNG contract prices, because LNG is an internationally traded commodity, and country-level customs data provides information on the annual LNG prices paid by an importer for LNG imported from different exporting countries. Using this data, it is possible to re-construct the contract price formulas the buyer and seller originally negotiated, albeit imperfectly. The customs data thus provides a useful check of the estimation methodology, since I can compare the contract prices predicted by the model with the contract prices that we can infer from the customs data.

Inferring contract price formulas from customs data I collected customs data from the Thomson Reuters Eikon terminal, covering three of the biggest LNG importers - Japan, Korea, and China. For each importer, this dataset includes monthly LNG import volumes from each exporting country, as well as monthly LNG prices (in USD/MMBtu) paid to each exporting country. The coverage is most extensive for Japan, with the dataset covering a nearly 20 year period from 1998 to 2018. For China and Korea, the dataset extends from 2009 - 2018, but there are gaps in the Chinese data, with prices missing for most of 2012 and 2013. All told, these data covers 22 different export-import country pair and a total of 4,522 exporting country-importing country-month observations.

The method I use for inferring contract price formulas from the customs data relies on the fact that LNG prices (particularly in Northeast Asia) are typically indexed to the price of oil through

[56] In the baseline specification, I use $\psi = 1.05$, but I find that different values of ψ lead to very similar approximations. The results are available upon request.

linear equations of the following form:[57]

$$p^c_{ijt} = a_{ij} + b_{ij} \sum_{s=0}^{T} w_s p^{oil}_{t-s} \tag{E5}$$

where p^c_{ijt} is the per-unit price of LNG that buyer j pays to seller i in period t (which is available from the customs data). $\sum_{s=0}^{T} w_s p^{oil}_{t-s}$ is the oil price index, which is a weighted average of current and past oil prices, with w_s the weight placed on the oil price in period $t-s$. T is the number of lags of the oil price included in the benchmark: for example, if $T = 3$, the benchmark oil price is a weighted average of the current oil price as well as the oil price in the last 3 months. a_{ij} and b_{ij} are the intercept and slope of the LNG pricing formula, and are negotiated at the time the seller and buyer sign the contract. We can see from the structure of equation (E5) that even though a_{ij} and b_{ij} are initially unknown (when the contract is signed), they can be eventually inferred by regressing the LNG price on lags of the oil price, once we have enough observations of p^c_{ijt} for the same exporter-importer pair i, j.

This approach of inferring contract pricing formulae from customs data is commonly used by industry analysts (see, for example, Flower and Liao, 2012). Agerton (2017) also relies on this approach in his empirical analysis of LNG contract pricing terms and the LNG price-oil price relationship. It is important to note, however, that there are some limitations of relying on customs data to recover contract pricing terms. First, customs data does not differentiate between long-term contracts and spot; as such, if importer j buys from exporter i under both long-term contracts and spot trade, the price p^c_{ijt} reported in the customs data will be a weighted average of the contract and spot price (where the weights are the shares of LNG imports accounted for by contracts and by spot). Second, an importer j and an exporter j may have multiple long-term agreements in force at any given time; in that case, the price p^c_{ijt} will be some weighted average of the individual contract prices, and it will generally not be possible to infer the individual contract prices. Third, some contracts (particularly those signed by Japanese buyers) use “S-curves" where they follow the linear pricing formula (E5) for most realizations of the oil price, but limit the dependence of the LNG price on oil prices (by using a smaller slope b_{ij}) when oil prices are either too high or too low. This is done as a way to reduce the exposure of the parties to oil market volatility.[58] Fourth, contract pricing terms may be renegotiated, in which case the pricing formula will change. Because of these issues, there is inevitably some noise in inferring contract pricing formulas from customs data (which is one reason why I do not use these inferred prices in structurally estimating the model).

[57]Oil indexation has been less dominant in other importing regions of the world; for example, when the US was a major LNG importer, many of its import contracts were reportedly indexed to domestic natural gas price indices (such as the Henry Hub)

[58]See, for example, Flower and Liao (2012), for more discussion on the use of S-curves in Asian LNG contracts.

Equation (E5) can be rewritten as follows, where $b_{ij}^{s} = b_{ij}w_s$:

$$p_{ijt}^{c} = a_{ij} + \sum_{s=0}^{t-T} b_{ij}^{s} p_{t-s}^{oil} \tag{E6}$$

I estimate LNG price formula coefficients a_{ij} and b_{ij}^{s} for each exporter-importer pair in the customs data by running linear regressions of p_{ijt}^{c} on lags of the Brent crude oil price. Because the number of lags actually used in the contract is unknown, I assume (following Agerton, 2017) that $T = 6$, so the regression includes both the current oil price, and lags of the crude oil price for up to 6 months in the regression. I restrict my analysis to the sub-sample of exporter-importer pairs that had at least one active long-term contract in the period.

Tables E10-E12 show the estimated LNG contract price formulas, starting with all of China's contracts, followed by Japan and Korea. For the most part the R^2 is quite high, but for some trade pairs (e.g. Papua New Guinea - China, Russia - Korea and Yemen - Korea) the R^2 is less than 0.5, suggesting that for these pairs, the customs data is noisy or the contracts used a different index.

Table E10: Estimates of Chinese LNG contract price formulas

Importer: China	Australia	Indonesia	Malaysia	PNG	Qatar
Intercept, a_{ij}	7.72***	8.83***	-0.100	4.77***	1.72***
	(0.44)	(0.52)	(1.21)	(0.43)	(0.46)
Slope coefficients, b_{ij}^{s}					
Current price, b_{ij}^{0}	-0.027	-0.069*	0.015	0.025	-0.0030
	(0.027)	(0.037)	(0.040)	(0.027)	(0.032)
1-month lag, b_{ij}^{1}	0.025	-0.019	-0.058	-0.038	0.030
	(0.043)	(0.058)	(0.057)	(0.042)	(0.051)
2-month lag, b_{ij}^{2}	-0.021	0.0078	0.083	0.026	-0.036
	(0.044)	(0.060)	(0.055)	(0.042)	(0.054)
3-month lag, b_{ij}^{3}	-0.011	-0.015	-0.023	-0.027	0.041
	(0.042)	(0.058)	(0.056)	(0.041)	(0.054)
4-month lag, b_{ij}^{4}	-0.022	0.019	0.030	0.049	0.027
	(0.042)	(0.059)	(0.056)	(0.041)	(0.051)
5-month lag, b_{ij}^{5}	0.029	0.0021	0.028	-0.076*	0.071
	(0.041)	(0.058)	(0.055)	(0.042)	(0.051)
6-month lag, b_{ij}^{6}	-0.013	0.041	0.083**	0.074***	0.016
	(0.025)	(0.036)	(0.033)	(0.025)	(0.032)
No. of observations	84	60	37	82	61
R^2	0.42	0.55	0.85	0.37	0.93

Compare model-predicted contract prices with contract prices inferred from customs data
Next, I compare the contract prices predicted by the structural model with the contracts prices inferred from customs data. The raw prices reported in the customs data cannot be directly compared

Table E11: Estimates of Japanese LNG contract price formulas

Importer: Japan	Australia	Brunei	Indonesia	Malaysia	Oman	Qatar	Russia	USA
Intercept, a_{ij}	1.11***	0.36*	0.73***	0.30*	3.22***	0.51***	1.22***	1.94***
	(0.16)	(0.21)	(0.21)	(0.17)	(0.27)	(0.14)	(0.34)	(0.28)
Slope coefficients, b^s_{ij}								
Current price, b^0_{ij}	0.027*	0.0022	0.0085	0.0020	-0.0012	0.0010	-0.049**	0.063**
	(0.016)	(0.020)	(0.020)	(0.016)	(0.020)	(0.014)	(0.020)	(0.029)
1-month lag, b^1_{ij}	-0.032	0.0014	0.046	-0.0041	-0.036	0.0062	0.040	-0.022
	(0.026)	(0.033)	(0.033)	(0.027)	(0.034)	(0.023)	(0.032)	(0.049)
2-month lag, b^2_{ij}	0.011	0.013	0.021	0.00065	0.032	-0.0020	-0.0014	-0.020
	(0.027)	(0.034)	(0.034)	(0.027)	(0.035)	(0.023)	(0.032)	(0.050)
3-month lag, b^3_{ij}	-0.0017	0.0028	0.0059	0.0017	0.017	0.025	0.0034	0.027
	(0.027)	(0.034)	(0.034)	(0.027)	(0.035)	(0.023)	(0.032)	(0.049)
4-month lag, b^4_{ij}	0.065**	0.023	0.020	0.073***	0.016	0.025	0.061*	0.0041
	(0.027)	(0.034)	(0.034)	(0.027)	(0.035)	(0.023)	(0.031)	(0.049)
5-month lag, b^5_{ij}	0.0063	0.0032	0.00091	0.025	0.025	0.028	0.039	0.024
	(0.026)	(0.033)	(0.033)	(0.027)	(0.034)	(0.023)	(0.031)	(0.050)
6-month lag, b^6_{ij}	0.048***	0.097***	0.037*	0.047***	0.022	0.060***	0.024	0.012
	(0.016)	(0.020)	(0.020)	(0.016)	(0.020)	(0.014)	(0.019)	(0.030)
No. of observations	240	247	247	247	206	247	112	171
R^2	0.92	0.91	0.90	0.94	0.71	0.95	0.91	0.66

Table E12: Estimates of South Korean LNG contract price formulas

Importer: Korea	Australia	Brunei	Egypt	Indonesia	Malaysia	Oman	Qatar	Russia	Yemen
Intercept, a_{ij}	1.76*	0.68**	0.52	1.25***	3.07***	1.14***	0.93**	5.68***	1.35
	(1.02)	(0.33)	(2.72)	(0.33)	(0.73)	(0.24)	(0.44)	(1.07)	(2.50)
Slope coefficients, b^s_{ij}									
Current price, b^0_{ij}	0.015	-0.0027	-0.038	-0.025	-0.082*	-0.0039	0.031	-0.082	-0.14*
	(0.062)	(0.021)	(0.084)	(0.020)	(0.044)	(0.014)	(0.026)	(0.067)	(0.071)
1-month lag, b^1_{ij}	0.045	0.011	-0.012	0.033	-0.027	-0.019	-0.047	0.040	0.040
	(0.099)	(0.033)	(0.16)	(0.031)	(0.068)	(0.023)	(0.041)	(0.11)	(0.11)
2-month lag, b^2_{ij}	-0.025	-0.039	0.20	0.083***	0.046	0.021	0.026	-0.058	0.024
	(0.11)	(0.033)	(0.14)	(0.031)	(0.068)	(0.022)	(0.041)	(0.10)	(0.12)
3-month lag, b^3_{ij}	-0.0044	0.054*	-0.16	-0.039	0.025	0.015	0.052	-0.023	0.060
	(0.095)	(0.032)	(0.16)	(0.031)	(0.068)	(0.023)	(0.041)	(0.099)	(0.12)
4-month lag, b^4_{ij}	0.069	0.030	0.048	0.013	0.097	0.046**	-0.055	0.13	-0.022
	(0.097)	(0.034)	(0.12)	(0.031)	(0.069)	(0.023)	(0.041)	(0.098)	(0.12)
5-month lag, b^5_{ij}	0.0029	0.050	0.14	0.034	0.0030	0.053**	0.033	0.085	0.010
	(0.092)	(0.035)	(0.13)	(0.031)	(0.068)	(0.023)	(0.041)	(0.098)	(0.12)
6-month lag, b^6_{ij}	0.012	0.049**	-0.037	0.012	0.031	0.040***	0.11***	-0.068	0.100
	(0.059)	(0.021)	(0.096)	(0.019)	(0.043)	(0.014)	(0.026)	(0.062)	(0.075)
No. of observations	81	92	28	110	109	110	110	101	64
R^2	0.52	0.95	0.65	0.90	0.63	0.97	0.90	0.19	0.36

to the model estimates, since the price series are truncated for most contracts.[59] As such, I use the estimated contract price formulas (reported in Tables E10-E12) together with information on oil prices to construct the *long-term average contract price* across the lifetime of each contract.[60] I then aggregate these contract-level average prices back to the exporter-importer level, using the contract quantities as the weights, to yield the average contract price p^c_{ij} that the exporter can be expected to pay the importer, based on the customs data.

The model estimates of the contract prices are constructed as follows. From the estimated structural model, we have estimates of T^c_{ij} (the lump-sum transfer) for every contract signed by importer j and exporter i, as well as the contract quantity they agree to (q^c_{ij}) and the contract duration. Together, these are used to calculate the the per-unit price for every contract. I again aggregate these contract-level average prices to the exporter-importer level, again using the contract quantities as the weights. This yields the model estimate of the average contract price that the exporter can be expected to pay the importer, $\hat{p}^c_{ij}$.

Figure E5 plots the model-predicted average contract price on the x-axis ($\hat{p}^c_{ij}$) and the average contract price inferred from customs data on the y-axis (p^c_{ij}). The two are fairly highly correlated, with a correlation coefficient of 0.45, suggesting the model provides reasonable estimates of contract prices. Table E13 shows summary statistics for the two sets of prices, aggregated to the exporter-importer level. The averages of the two prices are quite close to each other, and the model matches country-level average prices fairly well (for example, it is able to pick up the fact that China pays lower prices than Japan and Korea). The contract prices estimated from customs data tend to be more volatile than the model-predicted prices (with almost double the standard deviation), suggesting there are idiosyncratic differences between the prices of individual contracts that are difficult to explain via observed variables alone.

[59] Most contracts either began before the first year for which I have customs data, or end after the last year for which I have customs data. The truncation is especially severe for contracts signed by China and Korea, since I have only 10 years of customs data.

[60] This requires forming an expectation of future oil prices. I assume firms use a simple AR(1) model to forecast future prices.

Figure E5: Contract prices: model-predicted vs. customs data

Note: The figure plots average contract prices estimated from customs data (y-axis) against average contract prices that are predicted by the model (x-axis). Each dot represents an exporting country-importing country pair. The customs data is available for three importing countries (Japan, China, Korea) and their trading partners. See Appendix E.2 for details on how contract prices are computed from the customs data.

Table E13: Contract prices ($/MMBtu): model-predicted vs. customs data

Variable	Obs	Mean	Std. dev.	Min	Max
Contract price, customs data	22	12.09	3.79	3.32	17.87
Contract price, model-predicted	22	12.64	1.83	6.66	16.53
Contract price (China), customs data	5	10.14	6.49	3.32	17.87
Contract price (China), model-predicted	5	11.12	2.57	6.66	12.89
Contract price (Japan), customs data	8	12.04	1.78	8.68	13.92
Contract price (Japan), model-predicted	8	13.03	1.72	10.31	16.53
Contract price (Korea), customs data	9	13.21	3.18	8.61	17.03
Contract price (Korea), model-predicted	9	13.13	1.00	11.59	14.55

Additional References

CHEN, Y. AND D. Y. XU (2021): "A Structural Empirical Model of R&D, Firm Heterogeneity, and Industry Evolution," *Journal of Industrial Economics*.

FLOWER, A. AND J. LIAO (2012): "LNG Pricing in Asia," in *The pricing of internationally traded gas*, ed. by J. P. Stern, H. Wang, and T. Fan, Oxford: Oxford University Press Oxford Institute For Energy Studies, Oxford, chap. 11.

GERARDEN, T. (2017): "Demanding innovation: The impact of consumer subsidies on solar panel production costs," Working paper, Harvard University.

JEON, J. (2018): "Learning and Investment under Demand Uncertainty in Container Shipping," Department of Economics, The Boston University.

SWEETING, A. (2015): "A model of non-stationary dynamic price competition with an application to platform design," *Available at SSRN 2672027*.