

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

Nahim Bin Zahur *

May 11, 2025

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 significantly decrease investment, but allocative efficiency would improve. Policies aimed at eliminating contractual rigidities reduce investment by 30%, but raise welfare by 21%, with the industry becoming more nimble in responding to demand fluctuations.

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

*Department of Economics, Queen's University (email: nahim.zahur@queensu.ca). I am especially grateful to 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, Sylvain Chassang, Meredith Fowlie, Todd Gerarden, Matthew Gentzkow, Myrto Kalouptsi, Ryan Kellogg, Alex MacKay, Francesca Molinari, Mar Reguant, Ivan Rudik, Jorg Stoye, Lucy Xiaolu Wang, and seminar participants at the Bank of Canada, Cornell, Dallas Fed, Harvard, 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. Moreover, 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), pulp and paper (Darmouni et al., 2024), 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 can also trade on the spot. 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 this risk of under-investment. In addition, by providing predictable revenue streams, ex-ante long-term contracts can help sellers secure debt financing, reducing the cost of investment. On the other hand, such contracts may be inflexible in response to fluctuations in market conditions (compared to the alternative of trading on the spot), as it is costly to account for all possible contingencies when writing a contract (Masten and Crocker, 1985).

The trade-off between under-investment and contract inflexibility is particularly salient in the global liquefied natural gas (LNG) industry, one of the fastest-growing energy markets in the world. More than 70% of LNG is traded via long-term contracts that typically last for at least 20 years, with the rest traded on a spot market. LNG sellers must make significant investments (usually more than \$10 billion in cost) in liquefaction terminals that convert natural gas to LNG. Under-investment is a natural concern, both because of the difficulty of financing such large investments, and because 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. As such, sellers typically only invest *after* negotiating ex-ante long-term contracts with buyers. At the same time, LNG is subject to large demand fluctuations, such as the sharp increase in Japan's LNG demand following the Fukushima nuclear disaster in 2011 and the surge in European demand after the Russian invasion of Ukraine in 2022. Long-term 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, how does the use of long-term contracts affect investment and allocative efficiency? Second, are there welfare gains from regulating long-term contracts? In an effort to address the rigidity of long-term contracts, anti-trust authorities in

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

the European Union and Japan have introduced regulations limiting the use of resale restrictions. These regulations increase flexibility and reduce sellers' ability to exercise market power, thereby improving allocative efficiency. However, by shifting surplus from sellers to buyers, such regulations may unwittingly reduce sellers' incentives to invest in LNG terminals. The welfare impact of prohibiting resale restrictions thus hinges on the trade-off between under-investment and contract inflexibility that this paper seeks to quantify.

To answer these questions, this paper develops and estimates a structural model of the LNG industry that endogenizes both contracting and investment decisions, and embeds a micro-founded model of 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). Using the estimated model, I study the consequences of long-term contracts for investment and allocative efficiency, and assess the welfare effects of prohibiting resale restrictions in long-term contracts.

The model features spatially differentiated sellers and buyers.² Buyers have stochastic, unpredictable demand for LNG, and differ in their willingness-to-pay for contracted and spot LNG due to supply assurance motives (Bolton and Whinston, 1993), and transaction costs and search frictions (MacKay, 2022; Darmouni et al., 2024). Sellers incur fixed costs of investing in and financing liquefaction terminals. They face binding capacity constraints in the short-run, and incur 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. I assume the 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. In the third stage, after the seller has committed to the investment, sellers and buyers can negotiate *ex-post* contracts (once again via Nash-in-Nash bargaining).³ Finally, each year, demand shocks are realized and all agents participate in a global multi-regional spot market. Capacity-constrained sellers offer surplus LNG not already committed under long-term contracts and compete à la Cournot. 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 would under-invest if they were to only rely on *ex-post* contracts, both because they

²In the analysis, each exporting country is treated as a single seller, and each importing country as a single buyer. I discuss this assumption further in Section 2.

³Such *ex-post* contracts – where the negotiation takes place after the investment decision – may be valuable to buyers, since they do not require committing to trade as far in advance. In the LNG industry, *ex-post* contracts are (on average) signed 2.8 years before deliveries begin, compared to 5.3 years for *ex-ante* contracts.

cannot recoup the full marginal value of their investment when buyers have bargaining leverage, and because it is more costly to finance the investment without ex-ante contracts. Under-investment is mitigated by using ex-ante long-term contracts. The risk of under-investment (and the resulting incentive to sign larger ex-ante contracts) is larger the weaker the seller's bargaining leverage relative to the buyer. For instance, the risk of under-investment is greater when the seller is located far 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 this strengthens their outside option relative to the buyer.

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 depend on which of these two effects dominates.

The empirical analysis uses a rich dataset on capacity, long-term contracts, trade flows, and spot prices in the LNG industry, spanning the years 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 (an issue I return to below), I observe the duration, quantity and signature date for each contract, as well as the date of the "final investment decision" (i.e., the date when the seller commits to building the plant), start date, and capacity of each plant. The ability to observe the precise timing of contract signatures and investment decisions is a unique feature of the data and is crucial for identification of the structural model.

I use these data to document several descriptive facts consistent with the key mechanisms of the model. First, concerns about potential under-investment appear to be an important determinant of contracting behavior. Contracting primarily takes place before investment: on average, sellers commit around 60% of their capacity under ex-ante contracts before investing. To investigate this further, I measure 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). Consistent with the theory, ex-ante contracts are larger when sellers have weaker outside options, but smaller when buyers have weaker outside options.

Second, the upfront cost of investment is negatively correlated with the share of the project's capacity that the seller commits under ex-ante contracts. This is consistent with the mechanism that ex-ante contracts increase sellers' ability to obtain debt and lower the cost of financing.

Third, several pieces of evidence suggest that contractual rigidities bind in practice and inhibit the allocation of LNG. Sellers with a larger share of their capacity tied up in long-term contracts

are less responsive to short-run price differentials when allocating LNG across buyers. Moreover, resale restrictions in long-term contracts have led to the emergence of “reloads”, a very costly form of re-exporting where a contracted buyer takes physical delivery of the LNG cargo at their port (to comply with the contractual resale restriction), and then reloads it onto a different ship in order to re-sell. I find that LNG buyers that are more reliant on long-term contracts are more likely to engage in reloads, consistent with the presence of contractual rigidities.

I estimate the structural model in several steps. I first estimate demand curves for each buyer using demand shifters in other markets as instruments for the spot price. I then estimate seller production costs by exploiting the equilibrium conditions of the multi-regional 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 in turn determine their disagreement payoffs during contract negotiations.

Next, I estimate parameters characterizing the contracting and investment decisions: the investment cost, buyer preferences for contracting, and a Nash bargaining weight parameter that governs the distribution of surplus between sellers and buyers. 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 these outside options, identifies the bargaining weight. For instance, if the seller’s bargaining weight is high, changes in the seller’s outside option will have little impact on the negotiated price, and therefore have little effect on the equilibrium size of contracts and investments.⁴

I find that sellers face large sunk costs of investment, with an average export terminal estimated to cost \$9.8 billion to build. Ex-ante contracts lower the cost of financing, reducing the investment cost by 5.1% on average. The seller bargaining weight is estimated to be 0.63, and the hypothesis that sellers make take-it-or-leave-it offers is rejected by the data.⁵ Sellers’ incentives to invest are thus dampened by buyer bargaining power: the marginal benefit from investing would be 38% higher if sellers were able to fully capture 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 *ex-post* contracts: they are willing to pay a 7% premium to trade

⁴This strategy of leveraging the holdup effect to identify bargaining power shares similarities to [Bhattacharya \(2021\)](#), who uses information on ex-ante investments to recover a bargaining parameter.

⁵As 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.

using ex-post contracts rather than ex-ante contracts.

Using the estimated model, I carry out two sets of counter-factual exercises in order to answer the main research questions posed above. First, I evaluate the trade-off between under-investment and contract inflexibility. I begin by quantifying 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 sizable allocative efficiency gains, amounting to \$22 billion between 2006 and 2017. While removing long-term contracts worsens the deadweight loss from market power on the spot market, this is outweighed by the flexibility gains from freeing up seller capacity (that would otherwise be tied up in long-term contracts) in order to respond to demand fluctuations.

Next, I quantify the role of long-term contracts in mitigating under-investment. I find that if sellers were not able to sign long-term contracts with buyers, they would lower investment by 30.9%. The reduction in investment is primarily due to the inability to sign *ex-ante* contracts: if firms were able to sign ex-post but not ex-ante contracts, investment would still decrease by 27.8%. This reduction in investment is larger for sellers with less favorable outside options, such as those that are geographically more isolated or those making larger investments (who would find it more difficult to offload excess capacity on the spot market without contracts).

Second, I assess the long-term 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 from contract buyers; this weakens their outside option from trading on the spot market. As such, sellers reduce investment by 29.7%. Nonetheless, the removal of resale restrictions leads to a substantially more efficient allocation of LNG, generating welfare gains of \$326 billion (or 21.2%). Buyer surplus rises by 26.1%, but seller surplus declines by 4.5%, due to the decline in their bargaining leverage. This uneven distribution of gains may help explain why such policies have not been fully embraced by LNG industry participants.

To further investigate the allocative efficiency gains from increased flexibility, I study the effect of a shutdown of Russian natural gas exports to Europe, both in the baseline with rigid long-term contracts, and in the counter-factual where resale restrictions are removed. I find that rigid long-term contracts result in an inefficiently muted response to this demand shock, since some sellers bound by contracts do not reallocate LNG to European buyers experiencing higher demand; as a result, spot prices rise by 67%. When resale restrictions are lifted, more LNG is redirected to Europe, so that the demand shock only increases spot prices by 37%.

Contributions to the Literature: There is an extensive literature on long-term contracting and relationships (Williamson, 1975; Klein et al., 1978; Grossman and Hart, 1986; Joskow, 1987; Edlin and Reichelstein, 1996; Hubbard, 2001; Macchiavello and Morjaria, 2015; Macchiavello and Miquel-Florencia, 2017). This paper is most closely related to a recent strand of this literature that empirically examines – often through the lens of structural models – the value of long-term

contracts and relationships, and studies trade-offs across contractual forms (MacKay, 2022; Darmouni, Essig Aberg and Tolvanen, 2024; Harris and Nguyen, 2024; Garcia-Osipenko, Vreugdenhil and Zahur, 2025). In particular, Harris and Nguyen (2024) quantify trade-offs between long-term relationships and spot trade in the trucking industry, finding that relationships lead to thinner, less efficient spot markets. Garcia-Osipenko, Vreugdenhil and Zahur (2025) study the equilibrium effects of contract duration choices in containership leasing and show that contract rigidities generate capital misallocation across booms and busts.

My main contribution to this literature is to provide an empirical framework for studying the trade-off between under-investment and inflexibility that firms face when using long-term contracts. A novelty of the framework is that it endogenizes both contracting and investment decisions, while embedding a micro-founded model of the spot market. This enables me to quantify the extent to which long-term contracts mitigate under-investment, and to assess how contractual rigidities limit firms' responsiveness to demand shocks.

In complementary work that studies how LNG contracts operate on a week to week basis, Dhingra, Jia, Ottaviano, Sampson and Thomas (2023) find that contractual incompleteness enables sellers to better time their deliveries in response to demand shocks experienced by their contract buyers. In this paper, I show that these same contracts can simultaneously hinder the reallocation of LNG in response to demand shocks faced by buyers outside the contract.

This paper is also related to the literature on contract enforcement and hold-up risk (Nunn, 2007; Blouin and Macchiavello, 2019; Ryan, 2020; Bhattacharya, 2021; Ryan, 2021; Fabra and Llobet, 2025). The existing literature has emphasized inefficiencies from imperfect enforcement of contracts, such as under-investment arising from counterparty risk in long-term contracts (see, for example, Ryan, 2021 and Fabra and Llobet, 2025). By contrast, this paper considers a setting where long-term contracts are highly enforceable and therefore mitigate under-investment, yet can still generate inefficiencies from contract rigidity.

I contribute to a literature analyzing the interaction between long-term contracts and spot markets, mainly in the context of electricity. Building on the theoretical insights from Allaz and Vila (1993), a series of papers show that forward contracts mitigate the deadweight loss from seller market power in the spot market (Bushnell et al., 2008; Ito and Reguant, 2016). Unlike electricity forward contracts (which typically do not constrain allocations), long-term contracts for LNG and other commodities often include resale restrictions and other clauses that generate rigidity. I show that in the presence of such rigidities, contracts can potentially reduce allocative efficiency, despite their pro-competitive effect on the spot market (which I document exists for LNG as well).

The paper relates to a 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 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 (a common data challenge when studying business-to-business markets), but observes contracting and investment decisions that are functions of these negotiated prices.

Finally, the paper is also related to other papers that have studied long-term LNG contracts (Ruester, 2009; Hartley, 2015; Agerton, 2017). My paper builds on this literature by estimating a structural model of the LNG industry that can be used to quantify the impact of using long-term contracts and evaluate policies that reduce contract rigidity.

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 LNG tankers. LNG trade requires both the exporting and importing country to have 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.

The LNG industry has grown rapidly in recent years. 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 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.

Long-term contracts: The majority of LNG trade is carried out under long-term contracts signed between LNG suppliers and downstream buyers. A contract specifies the average annual contracted quantity, a start and end date, and a pricing formula indexing the LNG price to a selected benchmark (typically the oil price). Appendix Figure A2 plots the distribution of durations for long-term contracts (defined as contracts exceeding 4 years), showing that they typically last 15 years or longer.

In addition to these basic features, a contract generally specifies a “take-or-pay” quantity: a minimum quantity of LNG the buyer commits to paying for every year, whether or not they actually take delivery. Typical take-or-pay shares equal 90-100%, providing little flexibility to adjust quan-

⁶The closest antecedent to this strategy that I am aware of was developed by Bhattacharya (2021), who leverages the holdup effect to identify a bargaining parameter. A separate strand of the literature infers bargaining power in vertical supply chains where negotiated upstream prices are unobserved, by exploiting variation in firms’ relative bargaining positions that arises from heterogeneity in demand substitution patterns (Draganska et al., 2010; Barwick et al., 2025).

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

tity. Most contracts include *destination clauses*, which require that LNG cargoes be unloaded only within a pre-defined market (typically the buyer's home country), thus preventing buyers from easily re-selling LNG (IEA, 2013). Contracts may also include diversion clauses which specify how the two parties split the surplus if they decide to sell (or "divert") cargoes to third parties.

These resale restrictions (especially destination clauses) have been controversial, since they prevent buyers from easily engaging in arbitrage.⁸ Destination clauses have been challenged by anti-trust authorities in LNG importing countries, though with mixed success. The European Commission (EC) ruled destination clauses anti-competitive in 2003 (IEA, 2013), and has since then successfully negotiated their exclusion from LNG contracts signed with some (but not all) sellers (Talus, 2011). In 2017, the Japanese Fair Trade Commission (FTC) prohibited destination clauses in the majority of new LNG contracts (Harding and Sheppard, 2017). By the end of the study period in 2017, however, these policies had either not been fully implemented (in the case of Europe)⁹; or were not yet in place (in the case of Japan). As a result, the large majority of long-term contracts during this period have destination clauses.¹⁰

Re-negotiation of LNG contracts is rare: in my dataset (which I describe further below), 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, and the 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 contracts is costly both because of the negative reputational consequences and because of the risk of having to pay breach remedies.¹¹

Spot trade and reloads: The main alternative to long-term contracts is to trade LNG either on the spot market, or using short-term contracts.¹² Appendix Figure A3 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 faced with destination clauses can re-export LNG to other buyers by first unloading LNG from the seller in their terminal and re-loading the LNG onto a new tanker. However, such physical re-exports ("reloads") are a costly method of arbitrage, since they require the buyer to incur significantly higher shipping costs as well as additional operational costs (see Appendix

⁸Similar resale restrictions can be found in long-term contracts for coal (Joskow, 1985), wholesale electricity (Gavin and Ross, 2018) and pulp (Darmouni et al., 2024).

⁹A 2017 report by the EC stated that destination clauses had only been removed from "some contracts for LNG sales to Europe" (European Commission, 2017).

¹⁰For example, in Japan (by far the world's largest LNG importer), *all* contracted imports up until 2014 were covered by destination clauses; by 2017, the share was still as high as 96% (The Japan Fair Trade Commission, 2017).

¹¹For 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.

¹²A spot transaction involves delivery of a single LNG cargo from the seller's terminal to the buyer's terminal.

Figure A1). Not surprisingly, reloads only account for a small proportion of overall trade.¹³

LNG liquefaction and regasification plants: Liquefaction plants are significantly more costly than regasification plants: the capital cost of liquefaction plants has in recent years ranged from \$1.4 - \$2 billion per mtpa (million tonnes per annum) of capacity, compared to \$250 million per mtpa for regasification plants (OIES, 2017; Songhurst, 2018). For this reason, the binding constraint on trade tends to be available liquefaction capacity; capacity utilization for LNG exporters is high, equaling 92% on average (see Appendix Figure A4a). By contrast, there is a substantial amount of excess regasification capacity, with the average importer only operating at a capacity utilization of 40% (Appendix Figure A4b). In the model, I endogenize LNG sellers' investments in liquefaction plants, but abstract from buyers' decisions to invest in regasification plants.

Due to the high capital cost, financing liquefaction projects is challenging. LNG suppliers rely on equity and accrued cash flows, as well as *project finance*: a specialized form of debt finance commonly used for large infrastructure projects (see Appendix A.1 for details). Between 1985 and 2020, around 40% of the total capital for major liquefaction projects was funded using project finance (Baker, 2020). Project finance lenders (typically banks and export credit agencies) have traditionally required that suppliers secure long-term volume commitments from buyers sufficient to repay the debt. The ability to access cheaper project finance is therefore one reason for the widespread use of ex-ante long-term contracts; however, even LNG projects that are fully internally financed (e.g., by oil and gas majors) also rely heavily on ex-ante contracts (Baker, 2020).

Data

The empirical analysis utilizes historical data on the global LNG market, which I have collected from various industry sources; Appendix A.2 provides additional details.

Data on LNG contracts: Data on individual LNG contracts were collected from Bloomberg, the annual industry reports of the GIIGNL (The International Group of Liquefied Natural Gas Importers), and a dataset of natural gas contracts published by Neumann et al. (2015). In addition, for the vast majority of contracts, I have hand-collected press releases, company reports and newspaper articles announcing the signing of the contract. The eventual dataset consists of every long-term LNG contract signed 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 annual quantity, the signature year, the contract start and end year, the identity of the buyer and seller, the liquefaction plant used to fulfill the contract, and if the contract is an extension of an earlier, expiring contract. There are a few contract terms that I do not observe. The most important of these is the pricing formula, which is typically confidential and known only to the signatories of the contract. I also do not observe the "take-or-pay" share, or if the contract includes a destination or diversion clause.

¹³The share of reloads increased from 0.15% in 2008 to a peak of 2.69% in 2014, but decreased to 1% by 2017.

Data on liquefaction and regasification capacity: I obtain data on investment and capacity, for both liquefaction and regasification terminals, from the annual reports of the GIIGNL. The dataset includes, for every terminal operational from 2004 to 2017, the start-up year, nameplate capacity, the ownership structure and operator, the plant location and whether or not the plant is “greenfield” (i.e., built on a brand-new site, as opposed to being added to a site that already has an existing plant). 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 press releases and news articles.¹⁴ For the majority of plants, I also have data on the reported investment cost (including capitalised external finance costs) at the time of the FID.

The nameplate capacity of an LNG export plant may not accurately reflect the true level of capacity available in a given time period, since plants may be shut down for maintenance, and production may be disrupted by technical failures, sabotage, and disruptions in gas supply. Hence, I collect data on how many months each plant was operational in each year from 2004-17, and use this to construct the plant’s “effective capacity” each year, which equals its nameplate capacity multiplied by the fraction of months when the plant was operational (see Appendix A.2 for details).

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 aggregate these data to a country-year-quarter panel dataset recording LNG imports for each country, which I use for demand estimation. Second, I collect yearly LNG trade flows for each country pair from 2004 to 2017 from the GIIGNL’s annual reports, broken down into three types: long-term contracts; short-term contracts and spot trade; and reloads.¹⁵ I use this dataset when estimating the spot trade model. Finally, I obtain data on weekly LNG spot prices and shipping costs from several sources (see Appendix A.2).

Since trade flows are only observed at the country level, I treat countries as the decision-making agents in the empirical analysis: each exporting country is considered a single seller, and each importing country a single buyer. 16 out of 20 exporting countries have a single integrated firm in charge of all liquefaction facilities (usually the state-owned oil and gas company, such as Qatar Petroleum in Qatar); during the sample period, only Australia had more than two active sellers.¹⁶ LNG importing is somewhat more fragmented, but in contract negotiations, buyers from the same country very often negotiate as a consortium (Morikawa and Hashimoto, 2012).

Summary statistics Appendix Table A2 contains summary statistics on key variables used in the analysis. The average long-term contract is 17 years in duration, and is signed 3.6 years before

¹⁴The 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.

¹⁵Short-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.

¹⁶Since 2017, the market has become more fragmented due to entry by US LNG export facilities.

the start date of deliveries (Panel C). Export projects are generally very large in size (Panel D), with the average 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 & 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 traditional contracts with a fixed origin and 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 are significantly 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.

3 Descriptive Evidence

This section provides descriptive evidence suggesting that (i) concerns about under-investment are an important driver of contracting choices, and (ii) contractual rigidities contribute to short-run misallocation.

Ex-ante contracting and bargaining power Ex-ante contracts are widely used: as Figure 1 shows, contracting spikes during the months immediately before the investment decision. On average, a seller signs ex-ante contracts amounting to slightly over 60% of their capacity before they commit to the investment (Figure 2), with ex-post contracts only accounting for 23% of capacity. Appendix Table A4 shows that this pattern has remained stable over time, with sellers continuing

to rely heavily on ex-ante contracts for new investments. The growth of the spot market since 2004 has instead been driven largely by the expiry of legacy contracts.

Figure 1: Gap between contracting date and investment date





Figure 2: Cumulative share of capacity signed under long-term contracts



Note: Figure 1 shows the number of months between the contract signature date and the date when sellers make a final investment decision (FID). 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 2 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? By the time ex-post contracts are negotiated, the cost of investment is sunk and cannot directly influence the negotiated price. Depending on the bargaining position of the seller 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. This is plausible as there is only a limited pool of buyers to contract with at any given point in time (only 12 buyers negotiate long-term LNG contracts in an average year). There is also widespread heterogeneity in contract sizes signed within the same year, so that switching to a different buyer is likely to entail a large (and potentially costly) adjustment in the contract quantity.¹⁷ Finally, given the complexity of contract negotiations, switching to a different buyer is likely to take time and disrupt the seller's timetable – particularly important for debt-financed projects, where contracting and project financing must proceed in parallel.

Geography plays a crucial role in shaping the bargaining power of buyers relative to sellers, given the importance of shipping costs (on average 15% of the final price). If a seller is located close to a buyer they are negotiating with, but far from alternative buyers, switching partners would entail higher shipping costs, weakening the seller's bargaining position. Transaction costs theory (e.g., [Williamson, 1983](#)) predicts that in such a situation, sellers and buyers should sign larger ex-ante

¹⁷Likewise, buyers are constrained in how many sellers they can contract with, as documented in [Appendix A.5](#), so that neither side can easily switch trading partners.

contracts, as a way to mitigate under-investment. Conversely, when the buyer is located far away from potential alternative sellers, then the buyer's outside option is worse, reducing the incentive to contract ex-ante.

I now provide evidence that geography affects contracting behavior in the way suggested by the theory. I first compute a reduced-form proxy for how the bargaining leverage of each agent changes with geography, their *relative distance*, defined as their distance from alternative trading partners divided by the distance from their current trading partner (the one with whom they are negotiating the contract). The distance to alternative trading partners (the numerator) is measured by the 1st quartile of the distance between that agent and all potential trading partners.¹⁸ Intuitively, this measure captures the extent to which an agent's bargaining leverage is affected by geography: the larger the relative distance, the bigger the increase in shipping costs they have to incur in order to switch to a different trading partner, and therefore the weaker their bargaining position.¹⁹

Table 2: Geography and ex-ante contracting

Dependent variable	(1) Contract Quantity Share		(2) ln(Quantity)	
	Estimate	S.E.	Estimate	S.E.
Ex-ante*Relative distance, seller	0.17***	(0.044)	0.78**	(0.33)
Ex-ante*Relative distance, buyer	-0.15***	(0.037)	-0.82***	(0.27)
Relative distance, seller	-0.012	(0.017)	0.29**	(0.13)
Relative distance, buyer	0.011	(0.0091)	-0.042	(0.067)
Ex-ante contract	0.024	(0.035)	0.70***	(0.26)
Other Controls	Distance, Capacity, Extension, Time Trend, Greenfield			
N	337		337	
R ²	0.29		0.20	

Note: Each observation is a long-term contract. In specification (1), the dependent variable is the total contract quantity expressed as a share of plant capacity; in specification (2), it is the logarithm of the total contract quantity. The relative distance of an agent (which could be a seller or a buyer) is defined as their distance from alternative trading partners divided by the distance from their current trading partner; where the numerator, the distance to alternative trading partners, is defined as the 1st quartile of the distance between that agent and all potential trading partners. Other controls include the distance between the seller and buyer negotiating the contract, the logarithm of plant capacity, an indicator for greenfield plants, a dummy variable for whether the contract is an extension, and a time trend.

I then investigate how the use of ex-ante contracting varies with relative distance. In Column (1) of Table 2, I run a contract-level regression of the contract quantity (expressed as a share of the plant's capacity) on the relative distance of both the seller and the buyer, allowing the coefficients

¹⁸The idea behind using the 1st quartile is that, in the event of a contractual impasse, agents are more likely to switch to a trading partner located relatively close to them. The results are similar if I use instead the mean or median distance to other potential trading partners (Appendix A.6).

¹⁹This is similar to how Joskow (1987) classifies "mine-mouth" coal power plants as those located close to one coal supplier but far away from alternative coal suppliers; he finds that mine-mouth plants use longer ex-ante contracts.

to differ for ex-ante and ex-post contracts. Consistent with the theory, agents choose larger ex-ante contracts if the relative distance of the seller is higher or if the relative distance of the buyer is smaller (both of which decrease the seller's bargaining leverage). These effects are sizable in magnitude: a 10% increase in the seller's relative distance from alternative buyers raises the share of capacity contracted ex-ante by 0.17, while a 10% decrease in the buyer's relative distance from alternative sellers raises it by 0.15. Column (2) repeats the regression using the logarithm of the total quantity contracted as the dependent variable, finding very similar effects. Appendix A.6 shows the results are similar with alternative proxies for bargaining leverage, or controls for other determinants of quantity (such as past contracting). These regressions treat each contract as the unit of observation; aggregating the data to the investment level, I find that sellers who are located further away from alternative buyers negotiate ex-ante contracts amounting to a larger share of their total export capacity (Appendix Table A6), consistent with the contract-level evidence.

Ex-ante contracting and financing costs Another rationale for using ex-ante long-term contracts is that they may enable sellers to obtain more debt (and thereby lower the cost of financing the project), since securing ex-ante volume commitments from buyers lowers the risk profile of the investment project in the eyes of lenders. Consistent with this mechanism, Table 3 shows that the cost of investment for a project (which includes capitalised costs of external finance) is negatively correlated with the share of the project's capacity that is committed under ex-ante contracts.

Table 3: Ex-ante contracting associated with lower investment costs

Dependent variable	Average investment cost (USD bn/mtpa)	
	Estimate	S.E.
Share of ex-ante contracting	-0.792**	(0.367)
log(Capacity)	-0.178	(0.148)
Other Controls	Exporter Fixed Effects, Time Trend, Greenfield	
N	53	
R ²	0.61	

Note: Each observation is an investment project. The sample includes every investment whose final investment decision was made in 1995 or later, and for which data on investment cost is available. The investment cost includes the capitalised cost of external finance, as discussed in Appendix A.2. "Greenfield" investments are brand-new projects (as opposed to investments that are extensions of existing liquefaction projects).

Evidence for contract rigidity Next, I provide several pieces of evidence suggesting that contractual rigidities inhibit allocative efficiency.

Responsiveness of allocations to price differentials: First, I investigate whether sellers with a larger share of their capacity committed under long-term contracts are less responsive to short-term price differentials than sellers with mostly uncommitted capacity. To assess this, I estimate the

following regression using data on LNG trade flows:

$$s_{ijt} = \alpha_1 p_{ijt} + \alpha_2 (p_{ijt} \times \text{ShareContracted}_{it}) + \delta_{ij} + \delta_{it} + \varepsilon_{ijt} \quad (1)$$

s_{ijt} is the share of seller i 's output allocated to buyer j , and ranges from 0 to 1. p_{ijt} is the net price seller i receives from selling to buyer j , defined as the spot price paid by j minus the shipping cost. $\text{ShareContracted}_{it}$ is the proportion of seller i 's capacity committed under long-term contracts in period t ; it equals 0 for sellers who only sell on the spot, and 1 for sellers with fully committed capacity. I include seller-by-year fixed effects δ_{it} since my goal is to study how each seller's allocation of sales across buyers *within* a time period responds to cross-sectional price differences across buyers. I control for seller-buyer fixed effects δ_{ij} to account for time-invariant heterogeneity at the relationship level. Because the dependent variable s_{ijt} is bounded between 0 and 1, I estimate (1) as a fractional logit regression (Papke and Wooldridge, 1996). Since the net price p_{ijt} could be endogenous — as the seller's decision of how much to sell to buyer j might directly influence the spot price faced by buyer j — I instrument for the price using shipping costs and demand shifters, and employ a control function approach to incorporate these instruments.

The key parameters of interest are α_1 (which captures the baseline responsiveness of allocations to price) and α_2 (which captures how this responsiveness changes with the share of capacity committed). If contracts did not inhibit ex-post allocations (e.g., if take-or-pay clauses and/or resale restrictions did not actually bind in practice), allocations would be equally responsive to price differences regardless of the seller's contractual commitments, and α_2 would equal 0. However, I find that $\alpha_2 < 0$ (column (1) of Table 4), meaning that as a larger share of a seller's capacity is committed to long-term contracts, the seller becomes less responsive to price differences. In the limit where all of the seller's capacity is committed, the seller's allocations are completely unresponsive to price signals. I arrive at the same conclusion if I allow the price responsiveness to differ discretely for sellers with more than 75% of their capacity committed to long-term contracts (column (2)). Table A7 shows that these findings are robust to alternative transformations of the dependent variable (e.g., logarithm of s_{ijt} or inverse hyperbolic sine transformations).

Evidence from reloads: Second, I examine evidence from physical re-exports (or *reloads*). An example would be if a Spanish buyer received an LNG cargo from Qatar under a long-term contract, and re-loaded it onto a new tanker to be sold to a Japanese buyer. The very existence of such reloads is indicative of the presence of contractual rigidities (such as destination clauses), since shipping costs would be lowered and operational costs of reload avoided if the LNG cargo were instead sent directly from the seller to the eventual buyer: in the above example, if the LNG cargo were sent directly from Qatar to Japan, rather than taking a circuitous route from Qatar to Spain first and only then to Japan (see also Figure A1). Table 5 shows that buyers with a large share of their

Table 4: Allocations are less responsive to price differentials for sellers with a high share of capacity tied up under long-term contracts

Dependent Variable	Share of seller i 's output sold to buyer j			
	(1)		(2)	
	Estimate	S.E.	Estimate	S.E.
Net price, p_{ijt} (α_1)	1.52***	(0.16)	1.03***	(0.084)
Net price * Share Contracted (α_2)	-1.59***	(0.24)		
Net price * (Share Contracted > 0.75)			-0.96***	(0.14)
Fixed Effects	Seller-Year and Seller-Buyer FE			
N	3122		3122	
Pseudo- R^2	0.40		0.40	

Note: Each observation is an exporting country (seller) - importing country (buyer) - year. The dependent variable, s_{ijt} , is the share of seller i 's total LNG production that is sold to buyer j . The net price seller i receives from selling to buyer j , p_{ijt} is the difference between the spot price paid by buyer j , p_{jt} , and the the per-unit cost of shipping LNG from i to j , c_{ijt}^d is: $p_{ijt} = p_{jt} - c_{ijt}^d$. "Share Contracted" refers to the share of the seller's capacity that committed under long-term contracts. Since the dependent variable s_{ijt} is bounded between 0 and 1, I use a fractional logit regression (Papke and Wooldridge, 1996). To account for endogeneity of p_{ijt} , I follow a control function approach, where I first regress p_{ijt} on various IVs for the net price as well as seller-year and seller-buyer fixed effects, and then control for the residuals from that regression when estimating the fractional logit model. The IVs for p_{ijt} include shipping costs (c_{ijt}^d) and demand shifters that affect country j 's demand for LNG in period t (specifically, j 's electricity consumption from fossil fuels, coal price and the minimum temperature in country j).

Table 5: Buyers with a large share of imports from long-term contracts are more likely to reload

Dependent Variable Methodology	Indicator for positive reloads Logit		Resale quantity Tobit	
	(1)		(2)	
	Estimate	S.E.	Estimate	S.E.
Share of imports from long-term contracts	0.77***	(0.18)	0.13***	(0.031)
Net price from reloading, p_{jmt}	0.059***	(0.018)	0.0085***	(0.0023)
Positive net price ($p_{jmt} > 0$)	0.25**	(0.11)	0.051***	(0.019)
Fixed Effects	Importer and Year FE			
N	1991		1991	
Pseudo- R^2	0.22		0.39	

Note: Each observation is a reloading country (buyer j) - importing country (buyer m) - year (t). The key regressors are buyer j 's share of imports from long-term contracts in year t , and the net price j receives from reloading (i.e., physically re-exporting) LNG to buyer m , p_{jmt} . The latter is defined as $p_{jmt} = p_{mt} - p_{jt} - c_{jmt}^d$, and equals the difference in spot prices between the origin and destination market ($p_{mt} - p_{jt}$), adjusted for by the cost of shipping LNG from j to m (c_{jmt}^d). In (1), which is a logit regression, the dependent variable is an indicator for whether buyer j reloaded to buyer m in period t . In (2), which is a Tobit model, the dependent variable is the quantity of LNG buyer j reloaded to buyer m in period t .

imports from long-term contracts are significantly more likely to reload LNG (column (1)), and reload larger quantities of LNG (column (2)), consistent with the presence of contractual rigidities.

Other evidence: When LNG buyers experience major demand shocks (not uncommon in the industry), the extent to which sellers respond by re-directing LNG towards these buyers is often limited by contractual rigidities. For instance, when there was a large spike in Japan's LNG demand following the Fukushima nuclear disaster, Asian spot prices were \$5/MMBtu (nearly 50%) higher than European prices between 2011 and 2013 (Figure A7). This meant that sellers in the Middle East (such as Qatar and Algeria), who were roughly equidistant from Europe and Asia, would have received a \$5 spot premium by sending an LNG cargo to Asia than to Europe. However, during this period, these sellers continued to export LNG cargoes under long-term contracts to European buyers, who then re-loaded some of this LNG to Asia.²⁰ More recently, following the Russian invasion of Ukraine, US exporters in 2022 cited existing contractual commitments in Asia as a constraint on their ability to meet increased European demand for LNG.²¹

4 Model

4.1 Model overview and timing

Time is discrete and indexed by $t = 1, \dots, T$. Each period denotes a year. There are J buyers indexed by $j = 1, \dots, J$, and N sellers by $i = 1, \dots, N$. Sellers produce a homogeneous good (LNG). All agents are risk-neutral and have discount factor β . Buyers and sellers are spatially differentiated, with d_{ij} 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.²²

I model contracting and investment decisions for each seller i (and associated buyers) as a sequential, multi-stage game, as summarized in Figure 3.

Stage 1 of the game takes place before the seller has made any investment decision. Each seller-buyer pair negotiates an ex-ante contract: they Nash bargain 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 ex-ante contracts already negotiated in Stage 1, while anticipating both potential ex-post contracts they may sign in Stage 3 and the option to sell uncontracted capacity on the spot market. Ex-ante contracts lower investment

²⁰European reloads to Asia nearly quintupled between 2011 and 2013, compared to the 2008-2010 period. See Appendix A.8 for further discussion of the industry response to the Fukushima disaster.

²¹"World's Growing Thirst for American Gas Tests U.S. Ability to Meet Demand," *The Wall Street Journal*, July 2022.

²²For 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.

Figure 3: Stages of the game



costs by allowing sellers greater access to debt finance.

Stage 3 of the game occurs after the seller has committed to the investment. Each seller-buyer pair may now negotiate an ex-post contract, where they again Nash bargain over the contract quantity and a lump-sum transfer.²³

Finally, every year, all sellers and buyers participate in the LNG spot market (Stage 4). Demand shocks are realized. Sellers meet their contractual obligations, and sell uncommitted capacity on the annual spot market.

The model features the key economic mechanisms underlying contractual choices. First, *ex-ante* and *ex-post* contracting have differing implications for sellers' investment incentives. If sellers only rely on ex-post contracts (Stage 3), they cannot recoup the full marginal value of their investment if buyers have some bargaining power, and also face a higher cost of finance; both these channels lead to under-investment. Signing contracts ex-ante (in Stage 1) allows the seller and buyer to mitigate under-investment. However, both the magnitude of potential under-investment and the resulting incentive to sign ex-ante contracts depend critically on access to other trading partners via the spot market, with under-investment less of a concern for sellers with especially strong outside options.

Second, the use of long-run contracts has implications for allocative efficiency. Because contract quantities are fixed and determined 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. Each period, buyer demand shocks are realized. Sellers

²³These ex-post contracts may be negotiated either before or after the seller has completed construction of the plant.

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:

$$Q_{jt} = Q_d(p_{jt}, x_{jt}, \varepsilon_{jt}) \quad (2)$$

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. 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 spot price p_{jt} , since buyers are spatially differentiated. So in effect there are J different spot markets, each with its own price, though these prices are jointly determined in equilibrium since the same sellers can sell to each market.

Trade Flows: The quantity of LNG sold by seller i to buyer j equals $q_{ijt} = S_{ijt} + q_{ijt}^c$ where S_{ijt} denotes spot sales and q_{ijt}^c denotes contracted sales. Contracted sales q_{ijt}^c are pinned down by long-term contracts and cannot be adjusted in the short-term, whereas spot sales can. Seller i 's total production equals $q_{it} = \sum_j q_{ijt}$, while buyer j 's total imports equal $Q_{jt} = \sum_i q_{ijt}$.

Costs of production and sales: Sellers incur costs both in shipping the LNG to the buyers, and in producing LNG. Each unit of LNG costs c_{ijt}^d to ship from seller i to buyer j , where c_{ijt}^d is increasing in the distance d_{ij} between them and fluctuates by year.²⁴

Let $C(q_{it}, K_{it})$ denote the cost of production, where K_{it} is seller i 's effective capacity at time t . Capacity constraints are highly significant in this industry (as shown by Appendix Figure A4a). In the main analysis, I assume sellers face binding capacity constraints, so that q_{it} cannot exceed K_{it} .²⁵

Spot market equilibrium: 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 total short-run profit across all markets.²⁶

$$\begin{aligned} \{S_{ijt}\}_{j=1}^J = \operatorname{argmax}_{\{\hat{S}_{ijt}\}_{j=1}^J} & \left[\sum_{j=1}^J p_{jt}^* (\hat{S}_{ijt}, S_{-ijt}) \hat{S}_{ijt} - C(q_{it}, K_{it}) - \sum_{j=1}^J c_{ijt}^d q_{ijt} \right] \\ \text{s.t. } & q_{it} \leq K_{it} \end{aligned}$$

where p_{jt}^* is the market-clearing price in market j . The first-order condition satisfied by S_{ijt} (with equality if $S_{ijt} > 0$) is:

²⁴I assume shipping costs are the same for spot and contracted LNG. Consistent with this, Appendix Figure A5 shows that there is no systematic difference between short-term and long-term shipping rates; see Appendix A.4 for details.

²⁵In robustness checks, I consider other ways to model capacity constraints, finding similar results.

²⁶The short-run profit does not include revenue from contracted sales, since this is unaffected by spot market decisions and is effectively “sunk”. However the contract revenue is part of the seller’s long-run profits, and is recovered when I estimate the full contracting model (and accounted for in the subsequent welfare analysis).

$$\underbrace{p_{jt}^* + S_{ijt} \frac{\partial p_{jt}^*(S_{ijt}, S_{-ijt})}{\partial S_{ijt}}}_{\text{Marginal revenue of selling to market } j} - \underbrace{\left(\frac{\partial C(q_{it}, K_{it})}{\partial S_{ijt}} + \lambda_{it} + c_{ijt}^d \right)}_{\text{Marginal cost of selling to market } j} \leq 0 \quad (3)$$

The Lagrange multiplier λ_{it} can be interpreted as seller i 's “marginal profit” from increasing production: it equals 0 for unconstrained sellers, but may be positive for sellers operating at their capacity constraint.

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

Discussion: I assume that the quantity traded under each long-term contract cannot be adjusted once the contract has been signed: therefore, in the short-run, only spot trade flows can adjust to a demand shock. This is motivated by the widespread use of contractual clauses that inhibit flexibility (such as take-or-pay clauses and various resale restrictions), as well as the descriptive evidence of contract rigidity provided in Section 3.

While seller capacity constraints are a key feature of the model, I do not model buyer capacity constraints, since these almost never bind in the data (buyer utilization is less than 80% in over 95% of observations). A refinement of the model with capacity-constrained buyers yielded almost identical allocations as the baseline model, but is computationally much harder to solve.

Sellers are strategic and have market power on the spot market. A companion paper studying the LNG spot market provides evidence that sellers exercise market power (Zahur, 2023).

I assume buyers cannot engage in arbitrage, making it possible for sellers to practice spatial price discrimination. The most important barriers to arbitrage are destination clauses in long-term contracts; these prevent contracted buyers from engaging in the most cost-effective form of arbitrage, where they directly “divert” the LNG cargo from the seller to the terminal of the third party. Buyers faced with destination clauses can unload the cargo from their contract seller at their terminal and then reload the cargo onto a different ship, but re-exporting LNG in this fashion is costly because of the extra shipping costs involved (as shown in Appendix Figure A1) and additional costs incurred in reloading LNG. 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 (see 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 between seller i and buyer j consists of a start date, an end date, an annual contract quantity q_{ij}^c , 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 another.²⁷

Beliefs and expected payoffs: A key assumption I make is that demand shocks, ε_{jt} , are not serially correlated. 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 this assumption 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, so that demand shocks are potentially heteroskedastic: specifically, I assume $\varepsilon_{jt} \sim N(0, \sigma_j^2)$, where σ_j varies by j .

Let $\mathbf{q}_i^{c,1} = \{q_{ij}^c\}_{j \in \mathbf{B}_i^1}$ and $\mathbf{q}_i^{c,3} = \{q_{ij}^c\}_{j \in \mathbf{B}_i^3}$ denote vectors comprising all the contract quantities signed by seller i in Stage 1 and Stage 3 respectively. Let $\mathbf{Y}_{-i} = (\mathbf{q}_{-i}^c, \mathbf{K}_{-i})$ be a vector of the contracting and investment decisions made in projects operated by sellers other than i . Seller i and its contract buyers take $\mathbf{Y}_{-i}$ as given when choosing contracts and investment levels. Let t_i^3 denote the time period when Stage 3 contracts have been finalized. Seller i 's expected lifetime payoff, V_i^3 , can be written as the discounted sum of their expected profits in each future period:²⁸

$$V_i^3(\mathbf{q}_i^{c,1}, \mathbf{q}_i^{c,3}, K_i, \mathbf{Y}_{-i}) = \sum_{t=t_i^3}^{\infty} \beta^{t-t_i^3} \underbrace{\pi_{it}^s(\mathbf{q}_t^c, \mathbf{K}_t)}_{\text{Expected per-period variable profit}} \quad (4)$$

The buyer's expected payoff includes the discounted sum of their expected consumer surplus in each period, $\pi_{jt}^b(\mathbf{q}_t^c, \mathbf{K}_t)$. In addition, I allow for the possibility that buyers may be willing to pay a "premium" to purchase contracted LNG rather than spot LNG.²⁹ 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,

²⁷As discussed in Section 2, the large majority of LNG contracts during the sample period have destination clauses, and contracts are very rarely renegotiated or breached.

²⁸The superscript "3" is to indicate that these are their payoffs after Stage 3 is complete.

²⁹The premium may be negative for buyers who have a higher willingness-to-pay for spot LNG than contracted LNG.

2013). Long-term contracts may enable buyers to lower transaction costs and avoid trading frictions incurred on the spot market (MacKay, 2022; Dhingra et al., 2023; Darmouni et al., 2024). 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 shock to the marginal value of contracting between seller i and buyer j , which is publicly observable to all agents (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 :

$$\begin{aligned} W_j^3(q_i^{c,1}, q_i^{c,3}, K_i, Y_{-i}, \eta_{ij}^3) &= \sum_{t=t_i^3}^{\infty} \beta^{t-t_i^3} \underbrace{\pi_{jt}^b(q_t^c, K_t)}_{\text{Per-period consumer surplus}} + \underbrace{\omega_j^3(q_{ij}^{c,3}, \eta_{ij}^3)}_{\text{Contract premium, Stage 3}} \\ &= \underbrace{\tilde{W}_j^3(q_i^{c,1}, q_i^{c,3}, K_i, Y_{-i})}_{\text{Lifetime consumer surplus}} + \underbrace{\omega_j^3(q_{ij}^{c,3}, \eta_{ij}^3)}_{\text{Contract premium, Stage 3}} \end{aligned} \quad (5)$$

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 3), contracts are negotiated via Nash-in-Nash bargaining. Since lump-sum transfers are possible, each seller-buyer pair will (optimally) choose the contract quantity that maximizes their joint surplus. They will then negotiate the lump-sum transfer to divide the surplus from trading between the two parties (Chipty and Snyder, 1999).

Stage 3: ex-post contracting 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:

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

The FOC of the quantity problem is:

$$\frac{\partial V_i^3}{\partial q_{ij}^{c,3}} + \frac{\partial W_j^3}{\partial q_{ij}^{c,3}} = 0 \quad (6)$$

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

$$T_{ij}^{c,3} = \underbrace{\argmax_T \left(V_i^3(\mathbf{q}_i^{c,1}, \mathbf{q}_i^{c,3}, K_i, \mathbf{Y}_{-i}) - V_i^3(\mathbf{q}_i^{c,1}, \mathbf{q}_{i \setminus ij}^{c,3}, K_i, \mathbf{Y}_{-i}) + T \right)^\tau}_{\text{Seller's gains from trade in Stage 3}} \underbrace{\left(W_j^3(\mathbf{q}_i^{c,1}, \mathbf{q}_i^{c,3}, K_i, \mathbf{Y}_{-i}, \eta_{ij}^3) - W_j^3(\mathbf{q}_i^{c,1}, \mathbf{q}_{i \setminus ij}^{c,3}, K_i, \mathbf{Y}_{-i}) - T \right)^{1-\tau}}_{\text{Buyer's gains from trade in Stage 3}}, \forall j \in \mathbf{B}_i^3 \quad (7)$$

where τ is the bargaining weight of seller i when negotiating with buyer j .³⁰ 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 expected payoff for any buyer j at this point, W_j^2 , is their expected payoff at the end of Stage 3, W_j^3 , minus any transfers made to the seller.

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

$$W_j^2(\mathbf{q}_i^{c,1}, K_i, \mathbf{Y}_{-i}, \eta_{ij}^3) = W_j^3(\mathbf{q}_i^{c,1}, \mathbf{q}_i^{c,3}, K_i, \mathbf{Y}_{-i}, \eta_{ij}^3) - T_{ij}^{c,3}(\mathbf{q}_i^{c,1}, \mathbf{q}_i^{c,3}, K_i, \mathbf{Y}_{-i}, \eta_{ij}^3)$$

Stage 2: investment In Stage 2, the seller chooses how much capacity to build. Let $\Gamma_i(K_i, \mathbf{q}_i^{c,1}, \eta_i^2)$ denote the cost of the investment. η_i^2 is a publicly observable shock to seller i 's marginal cost of investing. I allow the investment cost to depend on ex-ante contracts $\mathbf{q}_i^{c,1}$, since sellers may be able to obtain cheaper debt (in the form of project finance) to finance the investment if they have already secured ex-ante volume commitments ([Hartley, 2015](#)).

³⁰I assume a single bargaining weight τ for the main analysis; in robustness checks, I allow τ to differ by seller and buyer.

The seller chooses K_i to maximize their net lifetime benefit from investing:

$$K_i^* = \operatorname{argmax}_{K_i} [V_i^2(\mathbf{q}_i^{c,1}, K_i, \mathbf{Y}_{-i}, \eta_{ij}^3) - \Gamma_i(K_i, \mathbf{q}_i^{c,1}, \eta_i^2)] \quad (8)$$

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

$$\frac{\partial V_i^2}{\partial K_i} - \frac{\partial \Gamma_i(K_i, \mathbf{q}_i^{c,1}, \eta_i^2)}{\partial K_i} = 0 \quad (9)$$

Stage 1: ex-ante contracting In Stage 1, just as with Stage 3, I allow the buyer to receive a premium from signing a contract of quantity $q_{ij}^{c,1}$ with seller i , to capture 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 and W_j^1 respectively denote seller i 's and buyer j 's lifetime expected payoffs as functions of contract quantities chosen in Stage 1. These payoffs equal:

$$\begin{aligned} V_i^1(\mathbf{q}_i^{c,1}, \mathbf{Y}_{-i}, \eta_i^2, \eta_{ij}^3) &= V_i^2(\mathbf{q}_i^{c,1}, K_i, \mathbf{Y}_{-i}, \eta_{ij}^3) - \Gamma_i(K_i, \mathbf{q}_i^{c,1}, \eta_i^2) \\ W_j^1(\mathbf{q}_i^{c,1}, \mathbf{Y}_{-i}, \eta_i^1, \eta_{ij}^2, \eta_{ij}^3) &= W_j^2(\mathbf{q}_i^{c,1}, K_i, \mathbf{Y}_{-i}, \eta_{ij}^3) + \omega_j^1(q_{ij}^{c,1}, \eta_{ij}^1) \end{aligned}$$

Equilibrium Quantities, Stage 1

Each seller-buyer pair chooses the contract quantity that maximizes their joint payoff given the choices of the other pairs. The equilibrium quantities are:

$$q_{ij}^{c,1} = \operatorname{argmax}_{q_{ij}} [V_i^1(\mathbf{q}_i^{c,1}, \mathbf{Y}_{-i}, \eta_i^2, \eta_{ij}^3) + W_j^1(\mathbf{q}_i^{c,1}, \mathbf{Y}_{-i}, \eta_{ij}^1, \eta_i^2, \eta_{ij}^3)], \forall j \in \mathbf{B}_i^1$$

The FOC is:

$$\frac{\partial V_i^1}{\partial q_{ij}^{c,1}} + \frac{\partial W_j^1}{\partial q_{ij}^{c,1}} = 0 \quad (10)$$

Equilibrium Transfers, Stage 1

If bargaining between i and j breaks down, they move on to the next stage of the game. Let $V_i^1(\mathbf{q}_{i \setminus ij}^{c,1}, \mathbf{Y}_{-i}, \eta_i^2, \eta_{ij}^3)$ and $W_j^1(\mathbf{q}_{i \setminus ij}^{c,1}, \mathbf{Y}_{-i}, \eta_i^2, \eta_{ij}^3)$ denote the disagreement payoffs of seller i and buyer j , respectively, when they negotiate with each another. Here $\mathbf{q}_{i \setminus ij}^{c,1}$ denotes the vector of Stage 1 contract quantities if we set $q_{ij}^{c,1}$ to 0 but maintain all other contracts in $\mathbf{q}_i^{c,1}$. The lump-sum

transfers $T_{ij}^{c,1}$ that seller i receives from buyer j , under Nash-in-Nash bargaining, are given by:

$$T_{ij}^{c,1} = \argmax_T \left(\underbrace{V_i^1(\mathbf{q}_i^{c,1}, \mathbf{Y}_{-i}, \eta_i^2, \eta_{ij}^3) - V_i^1(\mathbf{q}_{i,\backslash ij}^{c,1}, \mathbf{Y}_{-i}, \eta_i^2, \eta_{ij}^3)}_{\text{Seller's gains from trade in Stage 1}} + T \right)^\tau \underbrace{\left(W_j^1(\mathbf{q}_i^{c,1}, \mathbf{Y}_{-i}, \eta_{ij}^1, \eta_i^2, \eta_{ij}^3) - W_j^1(\mathbf{q}_{i,\backslash ij}^{c,1}, \mathbf{Y}_{-i}, \eta_i^2, \eta_{ij}^3) - T \right)^{1-\tau}}_{\text{Buyer's gains from trade in Stage 1}}, \forall j \in \mathbf{B}_i^1$$

4.4 Equilibrium properties of the contracting and investment game

Ex-ante contracting and investment: A key property of the model is that the seller will tend to *under-invest* if they do not use ex-ante contracts. Under-investment arises through two channels: a hold-up effect from ex-post bargaining between the seller and buyer, and an increase in the cost of financing the investment.

First, if the seller only signs contracts ex-post and has limited bargaining power, 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 is τ , 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 outside option: 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 potentially be socially beneficial, by strengthening the seller's outside option and reducing under-investment.

Second, if sellers only rely on ex-post contracts, they may face a higher cost of financing the project using debt (and therefore a higher investment cost). This is because lenders may view the investment as more risky if the seller has not secured ex-ante volume commitments. In the model, this is captured by letting the cost of investment Γ_i differ based on the size of ex-ante contracts.

Foreseeing potential under-investment, 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 in Stage 2 (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 incurred 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, taking into account the cost of the investment. Thus, ex-ante contracts mitigate (though they may not entirely eliminate) under-investment.³¹ It follows that the more severe the potential risk of under-investment (either because the seller has a lower bargaining weight τ or because the seller's outside option is

³¹Even with ex-ante contracting, there may be some under-investment relative to the optimum if ex-ante contracting involves opportunity costs, such as the need to commit to purchasing LNG many years in advance.

weak relative to the buyer), the larger the size of ex-ante contracts. Additionally, ex-ante contracts may also lower the cost of finance and therefore the investment cost. Supplementary Appendix [S.1.1](#) provides further details as well as numerical simulations illustrating these effects.

Mitigating the risk of under-investment is not the only reason to sign long-term contracts. Long-term contracts may be preferred to spot purchases due to supply assurance concerns, transaction costs of repeatedly buying on the spot, and risk aversion. Such motives (captured in the model through the contract premia) may explain why firms also use ex-post long-term contracts in addition to ex-ante contracts, as I explore further in the empirical analysis.

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 can reduce the ability of the industry to flexibly respond to demand shocks, especially in the presence of contractual resale restrictions (such as destination clauses). 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 sellers with a larger share of their output sold under contracts act more aggressively in the spot market, since they have less to lose by driving the spot price down (as the price they receive on the contracted portion of their output is unaffected by the spot price). Thus, it is an empirical question whether contracts improve or hurt allocative efficiency.

Long-term contracts also impose externalities on buyers and sellers who are not party to the contract. The intuition is simple: the two contracting parties maximize their bilateral surplus from trade, but do not internalize the impact (whether positive or negative) on other agents. A contract can impose negative externalities on “excluded” buyers because it commits a part of the seller’s output for the exclusive use of one buyer and reduces the quantity of LNG available to all other buyers on the spot market, who end up paying higher spot prices (especially during periods of high demand). Long-term contracts also impose externalities on excluded sellers, since sellers directly compete with each other on the spot market.³²

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.³³ 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

³²Supplementary Appendix [S.1.2](#) reports results from Monte Carlo simulations of the model that further explore the allocative efficiency consequences of rigid contracts, as well as contracting externalities.

³³This 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 (potentially with other trading partners) in Stage 3.

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 options. 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, given their rarity in practice (as discussed in Section 2).³⁴ By contrast, in settings where contract enforcement is weak, there is a significant risk that parties will breach contracts or opportunistically renegotiate contracts ex-post; then, under-investment can be significant even when ex-ante long-term contracts are used (Ryan, 2020, 2021; Fabra and Llobet, 2025).

In this model, firms guard against under-investment through ex-ante long-term contracts. An alternative organizational remedy is vertical integration (Grossman and Hart, 1986). Full vertical integration, however, is rare 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.

Finally, the model focuses on explaining the size distribution of ex-ante and ex-post long-term contracts, and takes their duration as given: empirical models of optimal contract duration have been developed in MacKay (2022) and Garcia-Osipenko et al. (2025).

5 Estimation

In this section, I describe how I estimate the model and discuss estimation results.

5.1 Estimation of Demand Curve

The demand curve (2) is parameterized as follows:

$$\ln(Q_{jt}) = \alpha - bp_{jt} + x_{jt}\theta_{dx} + \theta_j + \theta_{quarter} + \varepsilon_{jt} \quad (11)$$

I use quarterly panel data to estimate demand, and then aggregate demand to the annual level; to simplify the exposition, I use t to refer to an year-quarter in the discussion that follows (at the risk of abusing notation). In the baseline specification, x_{jt} includes electricity consumption from

³⁴The use of indexed (rather than fixed) prices further limits incentives to breach (Blouin and Macchiavello, 2019).

fossil fuels (in logarithms), the oil price, and the minimum temperature reached during period t (normalized to have zero mean and a standard deviation of 1). 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.³⁵ 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. 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 demand across countries (such as access to piped natural gas), and quarter fixed effects, $\theta_{quarter}$, to account for seasonal fluctuations.

I use three instruments for country j 's spot price p_{jt} , each of which shifts other countries' demand for LNG, while not directly affecting country j 's demand. The first is the average minimum temperature (normalized) in rival importers during the same period. When rival importers experience colder weather than usual, 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 and third instruments are rival electricity consumption from fossil fuels, and rival electricity generation from baseload sources (nuclear, renewables, hydroelectricity). If electricity demand increases in rival importers, or if their electricity generation from baseload sources falls, their demand for LNG rises, causing the LNG price paid by country j to increase.³⁶ These instruments are exogenous if country j 's idiosyncratic LNG demand shocks, after controlling for its own electricity demand, are uncorrelated with electricity demand or electricity generation shocks in other countries.

The demand estimates are shown in Appendix Table B8. The coefficient on the spot price is negative and implies a demand elasticity of around -0.83 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 10.1, so weak identification is unlikely to be a concern. Supplementary Appendix S.2.2 shows that the demand estimates are robust to alternative choices of control variables and instruments.

5.2 Estimation of Cost Parameters

Since shipping costs c_{ijt}^d are directly observed in the data, the only object to estimate is the production cost, or $C(q_{it}, K_{it})$. I assume that firms face a constant marginal cost of production, c , up until

³⁵The results are very similar if I use the country's total electricity consumption (from all energy sources) instead.

³⁶For instance, the Fukushima nuclear disaster in Japan caused Japan's electricity generation from baseload sources to decrease, generating increases in spot LNG prices paid by other countries.

they hit the binding capacity constraint (at which point their marginal cost becomes infinity).

$$C(q_{it}, K_{it}) = cq_{it} \text{ for } q_{it} \leq K_{it}$$

Marginal production costs are identified by variation in sellers' capacity utilization as demand changes. Note that if all sellers were fully capacity-constrained all the time, marginal costs would not be separately identified from the Lagrange multiplier on the capacity constraint (λ_{it}), but in the data there is plenty of variation in capacity utilization (Figure A4a).

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

The multi-regional Cournot spot market equilibrium is challenging to solve in the presence of binding capacity constraints. I adapt a fixed point algorithm developed by [Alsabab et al. \(2021\)](#) to numerically solve for the Cournot equilibrium (see Appendix B.1 for details). [Alsabab et al. \(2021\)](#) prove that there is a one-to-one mapping between the true Cournot game (where firms choose allocations S_{ijt}) and a reduced-form representation of this game (where firms instead choose their "marginal profits" from increasing production, λ_{it}). So instead of solving for the equilibrium spot allocations S_{ijt} directly, the algorithm instead solves for the equilibrium marginal profits λ_{it} via a fixed point routine; once these are known, the spot market allocation can be directly recovered.

Appendix Table B9 presents cost function estimates. The marginal cost of production, c , is estimated to be \$5.42/MMBtu, which is consistent with typical practitioner estimates.³⁷ The parsimonious model leads to a good fit of prices, spot exports and spot trade flows, and adding more covariates does not measurably increase the model fit.³⁸

5.3 Estimation of Contracting and Investment Model Parameters

The investment cost, and contract premium terms in Stages 3 and 1, are parametrized as follows:

$$\Gamma_i(K_i, \mathbf{q}_i^{c,1}, \eta_i^2) = \gamma_1 K_i + \frac{\gamma_2}{2} K_i^2 + \gamma_3 (K_i - \sum_{j \in \mathbf{B}_i^1} q_{ij}^{c,1}) + \eta_i^2 K_i \quad (12)$$

$$\omega_j^3(q_{ij}^{c,3}, \eta_{ij}^3) = \theta_3 q_{ij}^{c,3} + \frac{\kappa_3}{2} (q_{ij}^{c,3})^2 + \eta_{ij}^3 q_{ij}^{c,3} \quad (13)$$

$$\omega_j^1(q_{ij}^{c,1}, \eta_{ij}^1) = \theta_1 q_{ij}^{c,1} + \frac{\kappa_1}{2} (q_{ij}^{c,1})^2 + \eta_{ij}^1 q_{ij}^{c,1} \quad (14)$$

³⁷The marginal cost includes both the cost of obtaining raw natural gas (which can range from \$2-\$4/MMBtu), and operating expenses incurred by the liquefaction plant (typically \$1-\$2/MMBtu): see [Steuer \(2019\)](#).

³⁸In 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 with a quadratic or cubic cost function.

Thus, the average cost of investment (which I observe for most plants) equals:

$$\frac{\Gamma_i(K_i, \mathbf{q}_i^{c,1}, \eta_i^2)}{K_i} = \gamma_1 + \frac{\gamma_2}{2} K_i + \gamma_3 \frac{(K_i - \sum_{j \in \mathbf{B}_i^1} q_{ij}^{c,1})}{K_i} + \eta_i^2 \quad (15)$$

The investment cost is a quadratic function of investment K (with parameters γ_1 and γ_2), but in addition depends on the seller's *uncommitted capacity* ($K_i - \sum_{j \in \mathbf{B}_i^1} q_{ij}^{c,1}$): this is the total amount of capacity for which the seller does not have a contracted buyer at the time when they seek financing for the investment. Thus, the marginal cost of investment at $K = 0$ is γ_1 for a seller who commits all of their capacity under ex-ante contracts, but is $\gamma_1 + \gamma_3$ for a seller who only relies on ex-post contracts or the spot market. The parameter γ_3 therefore governs the extent to which a seller that choose not to use ex-ante contracts faces higher financing costs (e.g., because they are unable to secure as much debt, or have to pay higher interest rates on debt). Finally γ_2 determines how the marginal cost of investing varies with the level of investment.

The contract premium in both stages is assumed to be 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 Stages 1 and 3 because buyer preferences for contracts could depend on when the contract is signed. For example, θ_1 may be lower than θ_3 if buyers prefer not to sign contracts too far in advance when they have less information about their future demand.

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 , γ_2 and γ_3 (investment cost parameters) and τ (bargaining weight). Estimation utilizes the first-order conditions of each of the three stages of the game (equations (6), (9), and (10)); as well as equation (15) (which can be directly taken to the data, since I observe the average investment cost for each plant). I estimate the parameters using non-linear least squares.

Identification: The ex-post contract premium parameters, θ_3 and κ_3 , are identified directly from equation (6), which characterizes ex-post contract quantities.

The first-order condition characterizing investment (equation (9)) identifies the investment cost parameters γ_1 and γ_2 : intuitively, these parameters are pinned down by variation in investment as the seller's value of investing changes (for example, as the availability of buyers changes over time). The parameter capturing financing costs, γ_3 , is identified from equation (15) by exploiting variation in the (observed) average cost of investment as the share of capacity committed under ex-ante contracts changes: recall from Table 3 that plants with a greater share of ex-ante contracts have a lower average investment cost.

If contract prices were 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, I do not observe contract prices. Instead, I use equation (10) (which characterizes ex-ante contract quantities) to identify both τ , as well as the ex-ante contract premium parameters θ_1 and κ .

τ is identified from the relationship between ex-ante contract quantities and 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 mitigate under-investment – but the extent to which they will do so depends on τ . The greater τ is, the larger the share of the surplus the seller captures in ex-post negotiations, so the *less* sensitive ex-ante contract quantities are to the seller's disagreement payoff.³⁹ Variation in the buyer's disagreement payoff is likewise helpful for identifying τ : the greater τ is, the *more* sensitive 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 players 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.

With τ identified, the residual variation in ex-ante contract quantities helps identify θ_1 and κ_1 , the contract premium parameters in Stage 1. While I adopt a parsimonious quadratic specification for the contract premium in practice, this identification strategy would also extend to more flexible specifications of the Stage 1 contract premium as functions of observables, provided that these observables are not perfectly collinear with the disagreement payoffs of sellers and buyers.

Approximation of buyer and seller payoffs: In order to estimate the model, for each parameter guess, I 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 that constructing the exact payoffs is computationally very demanding, since it requires numerically solving for the spot market equilibrium every year while integrating out demand shocks.

Instead, I employ 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 $u_1, \dots, u_{L_s}$, and

³⁹This intuition is clearest in the polar cases where τ is either 1 or 0. For instance, when τ is 1, the buyer only needs to 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).

each buyer's per-period payoff by a set of L_b basis functions $\phi_1, \dots, \phi_{L_b}$:

$$\pi_{it}^s(\mathbf{q}_t^c, \mathbf{K}_t) \simeq \sum_{l=1}^{L_s} b_l^s u_l(\mathbf{q}_t^c, \mathbf{K}_t, \mathbf{x}_t); \quad \pi_{jt}^b(\mathbf{q}_t^c, \mathbf{K}_t) \simeq \sum_{l=1}^{L_b} b_l^b \phi_l(\mathbf{q}_t^c, \mathbf{K}_t, \mathbf{x}_t)$$

where b_l^s and b_l^b are unknown approximating parameters that need to be estimated. $\mathbf{x}_t$ is a vector of exogenous seller and buyer characteristics, and includes the distance between each seller and buyer, as well as buyer demand shifters x_{jt} .

The state space is a high-dimensional object, since it includes $\mathbf{q}_t^c$ (a vector of all contract quantities in period t) and a $\mathbf{K}_t$ (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 across 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 $\mathbf{q}_t^c$, $\mathbf{K}_t$, and $\boldsymbol{\varepsilon}_t$. 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. Supplementary Appendix [S.2.3](#) provides further details on the approximations, including the exact basis functions used.

Results: Table [6](#) shows the estimated parameters. I allow the coefficient governing the marginal cost of investment, γ_1 , to differ for greenfield projects (i.e., constructed on brand-new sites) and brownfield projects, finding the difference is positive but statistically insignificant. The quadratic coefficient, γ_2 , is estimated to be -0.16 (relative to a γ_1 of 31.95), suggesting modest economies of scale. γ_3 is positive (though imprecisely estimated) at 3.50, implying that for the median plant, ex-ante contracts reduce the average investment cost by 5.1%. This is consistent with a back-of-the-envelope calculation based on practitioner estimates of the financial benefits of debt, which suggests ex-ante contracting lowers investment costs by 6-7% (see Supplementary Appendix [S.2.5](#)).

Based on the parameter estimates, the average cost of building a 5 mtpa project (a median-sized project) is \$1.97 billion for every mtpa of capacity built. This is reasonably close to accounting estimates of the average cost, reported to be around \$1.4 - \$2 bn/mtpa ([Songhurst, 2018](#)).

The bargaining weight τ is estimated to equal 0.63, and I can reject the hypothesis that sellers have the ability to make take-it-or-leave-it offers to the buyer. An implication of this is that 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 38% on average. This, together with the financing channel, creates an incentive for sellers and buyers to sign larger ex-ante

Table 6: Contracting and investment parameter estimates

	Estimate	S.E.		Estimate	S.E.
Investment Cost Parameters			Contract Premium (ex-ante)		
γ_1	31.95	(7.77)	κ_1	0.0062	(0.0055)
$\gamma_1 \times \mathbb{1}\{\text{Greenfield}\}$	10.25	(7.77)	θ_1	0.26	(0.40)
γ_2	-0.16	(0.13)	Contract Premium (ex-post)		
γ_3	3.50	(4.66)	κ_3	-0.0012	(0.0022)
Bargaining Weight			θ_3	0.96	(0.06)
τ	0.63	(0.09)			
Number of contracts (Stage 1)	100		Number of investments	53	
Number of contracts (Stage 3)	194				

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

contracts to forestall under-investment.

The mean contract premium parameter in Stage 3 (θ_3) equals 0.96 (and is statistically significant), while κ_3 is negative and tiny. Together, these parameters imply that buyers are on average willing to pay a per-unit premium of \$0.95/MMBtu for LNG purchased under *ex-post* long-term contracts as opposed to spot purchases; the premium is about 11% of the average spot price of \$8.7/MMBtu. By contrast, the evidence for a contract premium in Stage 1 is weaker. The parameter θ_1 (the mean of the Stage 1 contract premium) is 0.26, and statistically indistinguishable from 0. While there are some increasing returns to contracting (since κ_1 is positive), the average per-unit premium for *ex-ante* contracts is only \$0.33/MMBtu, or 4% of the average spot price. These findings suggest that the main reason for using ex-ante contracts is to mitigate under-investment: if under-investment were not a concern, buyers would generally prefer ex-post over ex-ante contracts.

While long-term contracts are beneficial to the contracting parties, they also exert negative 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 would equal 0 on average. Instead, I find that the marginal externality is on average equal to -\$1.51/MMBtu, or about 17.5% of the average spot price.

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 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). Supplementary Appendix Figure S.5 shows that the model-predicted contract prices are quite similar to the contract prices inferred from customs data (with a correlation of 0.478), even though the latter data is not utilized in estimation.

Robustness checks: Appendix B.4 presents several robustness checks. Since the financing cost parameter γ_3 is imprecisely estimated, I estimate an alternative specification with γ_3 set to 0: this leads to a larger ex-ante contract premium, while all other parameters remain very robust (Table B11). Counter-factual analyses using this alternative specification yield results that are very similar to the baseline. Next, I explore specifications allowing heterogeneity in τ , finding little evidence that τ differs across sellers and buyers (Table B12). As such, I use the baseline specification with a single bargaining weight τ for the remainder of the analysis. Finally, I also find limited evidence of heterogeneity in the contract premium (Table B13).

6 Counter-factual Analysis

In this section, I use the model to explore the efficiency of long-term LNG contracts. In Section 6.1, I quantify the trade-off between under-investment and contract rigidity, by evaluating how the use of long-term contracts affects allocations and investment. In Section 6.2, I then study the long-run welfare impact of a policy that attempts to address contract rigidity by prohibiting the use of resale restrictions in long-term contracts.

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.⁴⁰ While this is a strong assumption (since it limits the richness of strategic interactions that are permitted in the counter-factual simulations), I find that the baseline simulations under this assumption yield a similar level of investment to what I see in the data.⁴¹

6.1 Decomposing the effects of long-term contracts on investment and allocative efficiency

Contracts and allocative efficiency: First, I quantify the allocative efficiency consequences of using long-term contracts. I simulate the spot market equilibrium, for a large number of realizations of demand shocks, both with and without long-term contracts. Since my goal here is to isolate how long-term contracts affects the short-run allocation of LNG, I hold the investment of every seller fixed in these counter-factuals.

Theory is ambiguous as to whether long-term contracts improve or worsen allocative efficiency. As discussed in Section 4.4, long-term contracts reduce the flexibility of sellers in meeting demand shocks, but decrease distortions from market power in the spot market. Table 7 shows that the inflexibility effect dominates, so that switching from long-term contracting to spot trade results in sizeable allocative efficiency gains. Short-run industry surplus between 2006 and 2017 (defined

⁴⁰In investment projects where a seller contracts with multiple buyers, I select the buyer with the largest contract.

⁴¹The capacity built in an average export project is 6.84 mtpa in the baseline simulations, versus 6.66 mtpa in the data.

as the sum of consumer surplus and producer profits) would increase by \$22 bn (2.4%) in the no-contracting regime relative to the benchmark regime (first row of Table 7).

These allocative efficiency gains are to some extent limited by the pro-competitive effects of long-term contracts on the spot market (due to the Allaz and Vila (1993) effect). If the spot market were perfectly competitive, eliminating long-term contracts would result in a significantly larger increase in surplus of \$42 bn (4.7%), as shown in the second row of Table 7.

These improvements in allocative efficiency largely arise because greater use of the spot market allows sellers to respond more efficiently to demand fluctuations. The last row of Table 7 shows that in the absence of any demand uncertainty, switching away from long-term contracts to spot trade would in fact *worsen* allocative efficiency (with short-run industry surplus reducing by 0.9%), due to increased market power distortions on the spot market.

Table 7: Allocative efficiency with and without long-term contracting, holding investment fixed

Industry short-run surplus (\$ bn)	Benchmark	No Contracting	% change
Baseline Assumptions	895	916	2.4%
Competitive Spot Market	901	943	4.7%
No Demand Uncertainty	823	815	-0.9%

Note: In the benchmark regime, sellers and buyers use long-term contracts (as observed in the data). In the “no-contracting” regime, no contracting is permitted. Investment is held fixed (at observed levels) in all simulations. For each counter-factual regime, I simulate 200 realizations of demand shocks for each buyer in each year from 2006 to 2017, and solve the spot market equilibrium for each realization. The table reports aggregate industry short-run surplus from 2006 to 2017 (discounted back to 2006), averaged across all simulations: here, the short-run surplus each year is defined as the sum of producer profits and consumer surplus in each year.

Finally, 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 little spare capacity on the spot market. As the contracted share of capacity decreases, though, these flexibility gains diminish, and the pro-competitive effect of long-term contracts begins to dominate. In Appendix C.3 I investigate the short-run impact of placing caps on the share of each seller’s capacity that can be contracted. I find that allocative efficiency is maximized by a cap of approximately 20%: beyond that point, further limits on long-term contracting lower allocative efficiency.

Contracts and investment: Next, I investigate how the ability to sign long-term contracts affects investment. I consider 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. In these “partial equilibrium” counter-factuals, I solve separately 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 contracting and investment incentives at the seller-buyer pair level. General equilibrium effects are explored in Section 6.2, when I evaluate the long-term impact of policies designed to reduce contract rigidities.

I find that that if sellers are not able to sign long-term contracts with buyers, they would lower investment by 30.9% on average (Table 8). 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 27.8%. These results indicate that ex-ante long-term contracts play an important role in mitigating under-investment, both through the bargaining channel and by reducing the cost of finance.

Table 8: Average capacity of project (mtpa), with restrictions on long-term contracting

	Average capacity (mtpa)	% change (average)	Correlation between % change and:	
			Distance from buyers	Capacity
Benchmark	6.84			
No contracting	4.72	-30.9%	-0.20	-0.54
No ex-ante contracting	4.94	-27.8%	-0.26	-0.55

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. All counter-factuals are partial equilibrium. The first two columns show the average capacity built by sellers, and (in italics) the percentage change relative to the benchmark regime. The last two columns report the correlation between the percentage change in capacity in the counter-factual regime, and (i) the distance from the seller to buyers (defined as the 1st quartile of the distance between the seller and all potential buyers, as discussed in Section 3) (ii) baseline plant capacity.

The last two columns of Table 8 show how the change in investment in each of these counter-factual regimes varies by seller characteristics. Sellers that are located further away from alternative buyers (and therefore have less favorable outside options from trading on the spot market) reduce investment more when they cannot use ex-ante long-term contracts. The reduction in investment from not being able to contract ex-ante is also bigger for sellers building larger plants, intuitively because such sellers have a weaker outside option when negotiating ex-post contracts with buyers – in the event of disagreement, they are left with a larger excess capacity to offload on the spot market.⁴² This makes them more reliant on ex-ante long-term contracts.

Summary: These counter-factuals quantify the central trade-off with using long-term contracts: they result in a more inefficient short-run allocation of LNG (in the presence of demand uncer-

⁴²Williamson (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.

tainty), but investment would be substantially lower if ex-ante long-term contracts were not used. The next section explores the regulatory implications of this trade-off.

6.2 Effects of banning resale restrictions

The inflexibility of long-term contracts is in part because of the widespread prevalence of resale restrictions (particularly destination clauses). 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 discuss a counter-factual that assesses the long-run consequences of the removal of all resale restrictions from LNG contracts.

In a world without contractually imposed resale restrictions, contracted buyers can re-sell LNG in response to unanticipated demand shocks elsewhere, increasing flexibility. They can also arbitrage away inter-regional price differentials through resales, weakening sellers' ability to engage in spatial price discrimination and thus reducing distortions from market power.

However, by reducing sellers' market power, the policy lowers their spot market profits while increasing buyer payoffs. This directly reduces sellers' incentives to invest. This also weakens the bargaining leverage of sellers relative to buyers in long-term contract negotiations, further worsening under-investment. The dilemma faced by regulators is therefore that a ban on resale restrictions may improve allocative efficiency but potentially reduce investment.⁴³

Modelling the spot market equilibrium when resale restrictions are removed: In the absence of resale restrictions, contracted buyers can now resell LNG, so sellers face the threat of arbitrage on the spot market. I assume this threat is strong enough that sellers are unable to price discriminate: they cannot charge different spot prices (net of shipping costs) to different buyers.⁴⁴ Therefore, for a given total quantity of LNG chosen by each seller, the *spatial allocation* of LNG is perfectly competitive, and unfettered by long-term contracts; sellers can only exercise market power through their choice of total production.

In a Cournot equilibrium with no resale restrictions, each seller i will simultaneously choose their total spot production, S_{it} ; their total quantity produced will therefore equal $q_{it} = S_{it} + \sum_j q_{ijt}^c$.

⁴³A similar policy trade-off arises in the regulation of the trade of pharmaceutical drugs: permitting parallel trade prevents drug manufacturers from price discriminating across buyers in different countries, but lowers the profits of manufacturers and reduce their incentives to innovate (Dubois and Sæthre, 2020).

⁴⁴For example, consider a single seller A selling to two buyers, 1 and 2 (as depicted in Figure A1). Suppose the spot prices paid by buyers 1 and 2 equal p_1 and p_2 , and the shipping cost from the seller to buyers 1 and 2 is c_1^d and c_2^d respectively. The net price the seller receives from selling to buyers 1 and 2 on the spot market thus equals $p_1 - c_1^d$ and $p_2 - c_2^d$. If the seller attempts to spatially price discriminate so that $p_2 - c_2^d$ exceeds $p_1 - c_1^d$, buyer 1 can buy LNG on the spot market (at a price of p_1) while asking the seller to "divert" some of their contracted cargoes to buyer 2 (at a price of p_2). Buyer 1 would need to compensate the seller for the difference in shipping costs $c_2^d - c_1^d$, but would still earn a net profit of $(p_2 - p_1) - (c_2^d - c_1^d)$, thereby defeating the seller's attempt to price discriminate.

Let $Q_t = \{q_{it}\}_{i=1}^N$ denote the vector of total quantities produced by each seller. Given any Q_t , let $A_t(Q_t) = \{q_{ijt}^*\}_{i=1, j=1}^{N, J}$ be the corresponding competitive spatial allocation of LNG. In this allocation, it is possible that a seller delivers less LNG to the buyer than contracted ($q_{ijt}^* < q_{ijt}^c$): if that is the case, then I assume the seller refunds buyer j for the shortfall in deliveries, with the refund price exactly equal to the prevailing spot price faced by buyer j .⁴⁵

Given that the spatial allocation of LNG is competitive, sellers will only end up selling LNG to those buyers from whom they receive the highest *net price* (i.e., price minus shipping costs). Let $p_{it}(S_{it}, S_{-it}) = \max_j (p_{jt} - c_{ijt}^d)$ denote this net price for seller i when they choose S_{it} given rival spot quantity choices $S_{-it} = \{S_{mt}\}_{m \neq i}$. The Cournot equilibrium conditions can then be written as:

$$\begin{aligned} S_{it}^* = \operatorname{argmax}_{S_{it}} & \left[p_{it}(S_{it}, S_{-it}) S_{it} - C(q_{it}, K_{it}) \right] \\ \text{s.t. } & q_{it} = S_{it} + \sum_j q_{ijt}^c \leq K_{it} \end{aligned} \quad (16)$$

Appendix C.1 describes a numerical algorithm for solving the resulting spot market equilibrium.

Thus, when resale restrictions are not present, long-term contracts no longer inhibit the *spatial* allocation of LNG. However, the larger the total contracted quantity of a seller, the weaker their incentive to lower their production as a way to drive up the spot price: so that the Allaz and Vila (1993) pro-competitive effect of contracts will still be present.

Modelling the long-term industry equilibrium when resale restrictions are removed: I assume that resale restrictions are prohibited in all long-term contracts (both existing and new), starting from the year 2012. This policy comes as an one-time, unanticipated shock to the industry. I simulate the new industry equilibrium by solving the multi-stage game of contracting and investment, using updated seller and buyer payoff functions that are derived based on the spot market equilibrium without resale restrictions. Unlike the partial equilibrium counter-factuals considered in Section 6.1, I now account for general equilibrium considerations. This means that each seller and their contracted buyers, when making investment and contracting decisions, account for how other sellers and buyers will adjust their investment and contracting choices in response to the removal of resale restrictions. Solving the full industry equilibrium is computationally involved and the Gauss-Jacobi fixed point algorithm for doing so is described in Appendix C.2.

Findings: I find that removing resale restrictions in long-term contracts would result in a large *decrease* in investment, by 29.7% on average (Table 9). As such, total capacity in the industry decreases by 9.1% in the long-run. The decline in investment is driven by a combination of two effects. First, the removal of resale restrictions reduces sellers' market power on the spot market,

⁴⁵This 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.

so the average spot price decreases by 6%: this directly reduces sellers payoffs from spot sales. Second, the policy worsens the bargaining leverage of sellers when negotiating long-term contracts with buyers, since the seller has a weaker outside option relative to the buyer (as evidenced by spot prices being lower). As a result, the average contract price decreases by 3.4%. 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 share of new capacity that is contracted increasing from 89% to 92%: it is less costly to sign large contracts in this environment, since any excess quantity contracted can always be re-sold. Overall, though, sellers still have weaker incentives to invest, resulting in the large decrease in investment.

Table 9: Investment and welfare impact of prohibiting resale restrictions

	Benchmark	No resale restrictions	% change
Welfare (\$ bn)	1,534	1,860	21.2%
Seller surplus (\$ bn)	244	233	-4.5%
Buyer surplus (\$ bn)	1,289	1,626	26.1%
Long-run capacity (mtpa)	445	404	-9.1%
Average capacity of new plant (mtpa)	7.15	5.03	-29.7%
Share of new capacity contracted	88.6%	92.0%	3.9%
Average contract price (\$/MMBtu)	7.92	7.66	-3.4%
Average spot price (\$/MMBtu)	9.42	8.86	-6.0%

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 long-run industry capacity (the total capacity reached by the year 2045). The fifth row shows the average capacity of “new plants”, defined as those plants where the investment decision is made after 2012, in million tonnes per annum (mtpa); while the sixth row shows the share of new capacity that is contracted. The final two rows show the average contract price and the average spot price (in \$/MMBtu) following the introduction of the policy in 2012.

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. Thus, despite the reduction in investment, there are sizable welfare gains from removing resale restrictions, with welfare increasing by around \$326 bn (or 21.2%).⁴⁶ As I explore further below, one source of these welfare gains is that the industry becomes more nimble in responding to major demand shocks.

The welfare gains from removing resale restrictions are not enjoyed by both buyers and sellers alike, though: while buyer surplus increases by 26.1%, seller surplus declines by 4.5%. The

⁴⁶This does not account for the fact that some sellers made investment decisions prior to the (unexpected) removal of resale restrictions: had they anticipated the policy change, they might have opted to invest less. Therefore, to bound the potential welfare gains from removing resale restrictions, I also simulate the evolution of the industry assuming that both new *and* existing plants reduce capacity by 29.7%, finding that welfare would still increase by 15.1%.

uneven division of gains between buyers and sellers helps explain why resale restrictions are still widely used in LNG contracts, despite their inefficiency. Unsurprisingly, LNG exporters have 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.⁴⁷ 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.

Implications for responses to demand shocks: To gain further insight into the allocative efficiency benefits of relaxing resale restrictions and increasing contract flexibility, I consider the impact of a 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, both in the baseline regime that has rigid long-term contracts, and the counter-factual regime where resale restrictions are removed.⁴⁸

Table 10: Impact of hypothetical shutdown of Russian natural gas exports: with and without resale restrictions

	Benchmark	No Resale Restrictions
Δ European LNG imports (mt)	49.0	58.9
<i>Δ Imports from sellers with low contract share</i>	22.8	21.2
<i>Δ Imports from sellers with high contract share</i>	26.1	37.7
Δ European Average Spot price (\$/MMBtu)	4.4	2.4
Δ Global Average Spot Price (\$/MMBtu)	3.5	2.6

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 is banned in 2012, so that from then on all buyers under long-term contracts can freely re-sell LNG. The table reports the impact of a temporary shutdown of Russian natural gas exports to Europe on allocations and spot prices under both these regimes; I assume that the shutdown of Russian natural gas exports cause demand for natural gas to increase by 100 bcm (73 mt) under fixed prices. “ Δ European LNG imports” (first row) shows the increase in European LNG imports due to the shutdown of Russian pipeline exports under the two different contracting regimes; the second and third two rows decompose this increase in imports into imports coming from sellers with less than 50% of their capacity committed under long-term contracts (“low contract share”), and those with at least 50% their capacity committed under long-term contracts (“high contract share”). The final two rows show how the shutdown affects both the sales-weighted average spot price paid by European buyers, as well as the sales-weighted global average spot price.

I find the industry responds more efficiently to the demand shock if resale restrictions were not present in long-term contracts (Table 10). In the event of a Russian natural gas export shutdown,

⁴⁷For example, destination clauses were a key sticking point in the LNG contract negotiations between Germany and Qatar in 2022. See <https://www.reuters.com/business/energy/exclusive-germany-qatar-odds-over-terms-talks-lng-supply-deal-sources-2022-05-09/>.

⁴⁸In Appendix C.4, I provide additional details and also carry out a similar exercise to study the LNG industry response to the Fukushima nuclear disaster in Japan in 2011.

European LNG imports rise by 49.0 mt annually under the baseline regime with rigid long-term contracts. But without resale restrictions, European imports would increase by 58.9 mt – 20% more – even though there is less liquefaction capacity available. As such, the demand shock raises spot prices in Europe by \$4.4/MMBtu (or 67%) in the baseline regime with rigid long-term contracts, but by only \$2.4/MMBtu (or 37%) when resale restrictions are absent.

This difference arises because rigid long-term contracts result in an inefficiently muted response to the demand shock: some sellers bound by long-term contracts do not re-allocate LNG to buyers in Europe experiencing increased demand. Consistent with this mechanism, Table 10 shows that the additional European imports (when resale restrictions are removed) largely come from sellers with a high share of their capacity committed under long-term contracts.⁴⁹

7 Conclusion

In markets where firms must make large sunk cost investments for trade to occur, ex-post bargaining can reduce the surplus they capture from the investment, leading to under-investment. Firms may also under-invest because of the difficulty of financing large investments. Ex-ante long-term contracts help mitigate under-investment from both these channels. However, contracting rigidities can inhibit the ability of firms to respond to demand shocks. In this paper, I quantify this trade-off and assess its policy implications in the context of the global LNG industry. My main contribution is to develop and estimate a novel empirical framework that endogenizes both contracting and investment decisions, while embedding a micro-founded model of the spot market.

Using the model, I find that without long-term contracts, sellers would significantly decrease investment, but allocative efficiency would improve. Policies that reduce contractual rigidities by prohibiting resale restrictions reduce investment by 30%, but lead to substantial welfare gains of 21%, largely driven by sellers becoming more responsive to large demand shocks. More broadly, these results underscore that the trade-off between under-investment and contract rigidity is central to understanding how contractual arrangements in many business-to-business markets are structured, and whether and how they should be regulated.

References

- Agerton, Mark**, “Global LNG pricing terms and revisions: An empirical analysis,” *The Energy Journal*, 2017, 38 (1).
- Allaz, Blaise and Jean-Luc Vila**, “Cournot Competition, Forward Markets and Efficiency,” *Journal of Economic Theory*, 1993, 59 (1), 1–16.

⁴⁹This tallies with the descriptive evidence presented earlier in Table 4, which showed that sellers with a high share of their capacity contracted are less responsive to short-run price differentials; unsurprisingly, these are the sellers that adjust allocations the most when contract rigidities are eliminated.

- Alsabah, Humoud, Benjamin Bernard, Agostino Capponi, Garud Iyengar, and Jay Sethuraman**, “Multiregional oligopoly with capacity constraints,” *Management Science*, 2021, 67 (8), 4789–4808.
- Ason, Agnieszka**, “Price reviews and arbitrations in Asian LNG markets,” 2019. Oxford Institute for Energy Studies: NG 144.
- Baker, R**, “LNG Finance-will lenders accommodate the changing environment?,” *Oxford Institute for Energy Studies*, 2020, pp. 1–29.
- Barwick, Panle Jia and Parag Pathak**, “The Cost of Free Entry: An Empirical Analysis of Real Estate Agents in Greater Boston,” *RAND Journal of Economics*, 2015, 46 (1), 103–145.
- , **Hyuk-Soo Kwon, Shanjun Li, and Nahim Bin Zahur**, “Drive Down the Cost: Learning by Doing and Government Policies in the Global EV Battery Industry,” 2025. National Bureau of Economic Research Working Paper.
- Benkard, C Lanier, Przemyslaw Jeziorski, and Gabriel Y Weintraub**, “Oblivious equilibrium for concentrated industries,” *The RAND Journal of Economics*, 2015, 46 (4), 671–708.
- Besanko, David and Ulrich Doraszelski**, “Capacity dynamics and endogenous asymmetries in firm size,” *RAND Journal of Economics*, 2004, pp. 23–49.
- Bhattacharya, Vivek**, “An empirical model of R&D procurement contests: An analysis of the DOD SBIR program,” *Econometrica*, 2021, 89 (5), 2189–2224.
- Blouin, Arthur and Rocco Macchiavello**, “Strategic default in the international coffee market,” *The Quarterly Journal of Economics*, 2019, 134 (2), 895–951.
- Bolton, Patrick and Michael D Whinston**, “Incomplete contracts, vertical integration, and supply assurance,” *The Review of Economic Studies*, 1993, 60 (1), 121–148.
- Bushnell, James B., Erin T. Mansur, and Celeste Saravia**, “Vertical Arrangements, Market Structure, and Competition: An Analysis of Restructured US Electricity Markets,” *American Economic Review*, 2008, 98 (1), 237–266.
- Chipty, Tasneem and Christopher M Snyder**, “The role of firm size in bilateral bargaining: A study of the cable television industry,” *Review of Economics and Statistics*, 1999, 81 (2), 326–340.
- Crawford, Gregory S and Ali Yurukoglu**, “The welfare effects of bundling in multichannel television markets,” *American Economic Review*, 2012, 102 (2), 643–85.
- Darmouni, Olivier, Simon Essig Aberg, and Juha Tolvanen**, “Pulp Friction: The Value of Quantity Contracts in Decentralized Markets,” 2024. SSRN working paper 3919592.
- Dhingra, Swati, Ningyuan Jia, Gianmarco Ottaviano, Thomas Sampson, and Catherine Thomas**, “Global Supply Chain Flexibility: The Role of Contracts in LNG,” 2023. Working Paper.
- Draganska, Michaela, Daniel Klapper, and Sofia B. Villas-Boas**, “A Larger Slice or a Larger Pie? An Empirical Investigation of Bargaining Power in the Distribution Channel,” *Marketing Science*, 2010, 29 (1), 57–74.
- Dubois, Pierre and Morten Sæthre**, “On the effect of parallel trade on manufacturers’ and retailers’ profits in the pharmaceutical sector,” *Econometrica*, 2020, 88 (6), 2503–2545.
- Edlin, Aaron and Stefan Reichelstein**, “Holdups, Standard Breach Remedies, and Optimal Investment,” *American Economic Review*, 1996, 86 (3), 478–501.
- European Commission**, “Follow-up study to the LNG and storage strategy,” 2017.
- Fabra, Natalia and Gerard Llobet**, “The Costs of Counterparty Risk in Long-Term Contracts,” 2025. Working Paper.

- Garcia-Osipenko, Maria, Nicholas Vreugdenhil, and Nahim Bin Zahur**, “Endogenous Rigidities and Capital Misallocation: Evidence from Containerships,” 2025. Working Paper.
- Gavin, Sebnem and Thomas W Ross**, “Long-term contracts as barriers to entry with differentiated products,” *International Journal of Industrial Organization*, 2018, 59, 514–537.
- Gowrisankaran, Gautam, Aviv Nevo, and Robert Town**, “Mergers when prices are negotiated: Evidence from the hospital industry,” *American Economic Review*, 2015, 105 (1), 172–203.
- Grennan, Matthew**, “Price discrimination and bargaining: Empirical evidence from medical devices,” *American Economic Review*, 2013, 103 (1), 145–177.
- Grossman, Sanford J and Oliver D Hart**, “The costs and benefits of ownership: A theory of vertical and lateral integration,” *Journal of Political Economy*, 1986, 94 (4), 691–719.
- Harding, Robin and David Sheppard**, “Japan outlaws restrictions on resale of LNG cargoes,” *Financial Times*, 2017.
- Harris, Adam and Thi Mai Anh Nguyen**, “Long-term Relationships and the Spot Market: Evidence from US Trucking,” 2024. Working Paper.
- Hartley, Peter**, “The Future of Long-term LNG Contracts,” *The Energy Journal*, 2015, 36 (3).
- Ho, Kate and Robin S Lee**, “Insurer competition in health care markets,” *Econometrica*, 2017, 85 (2), 379–417.
- and —, “Equilibrium provider networks: Bargaining and exclusion in health care markets,” *American Economic Review*, 2019, 109 (2), 473–522.
- Horn, Henrick and Asher Wolinsky**, “Bilateral monopolies and incentives for merger,” *RAND Journal of Economics*, 1988, pp. 408–419.
- Hubbard, Thomas N**, “Contractual form and market thickness in trucking,” *RAND Journal of Economics*, 2001, pp. 369–386.
- IEA**, *Developing a Natural Gas Trading Hub in Asia: Obstacles and Opportunities*, Paris: OECD/IEA, 2013.
- Ito, Koichiro and Mar Reguant**, “Sequential Markets, Market Power, and Arbitrage,” *American Economic Review*, 2016, 106 (7), 1921–1957.
- Joskow, Paul**, “Vertical integration and long-term contracts: The case of coal burning electric generation plants,” *Journal of Law, Economics and Organization*, 1985, 1 (1), 33–80.
- , “Contract Duration and Relationship-Specific Investments: Empirical Evidence from Coal Markets,” *American Economic Review*, 1987, 77 (1), 168–185.
- Klein, Benjamin, Robert G. Crawford, and Armen A. Alchian**, “Vertical integration, appropriate rents, and the competitive contracting process,” *The Journal of Law and Economics*, 1978, 21 (2), 297–326.
- Macchiavello, Rocco and Ameet Morjaria**, “The value of relationships: evidence from a supply shock to Kenyan rose exports,” *American Economic Review*, 2015, 105 (9), 2911–45.
- and **Josepa Miquel-Florensa**, “Vertical integration and relational contracts: Evidence from the Costa Rica coffee chain,” 2017. CEPR Discussion Paper No. DP11874.
- MacKay, Alexander**, “Contract duration and the costs of market transactions,” *American Economic Journal: Microeconomics*, 2022, 14 (3), 164–212.
- Masten, Scott E. and Keith J. Crocker**, “Efficient adaptation in long-term contracts: Take-or-pay provisions for natural gas,” *American Economic Review*, 1985, 75 (5), 1083–1093.
- Morikawa, Tetsuo and Hiroshi Hashimoto**, “Japan’s new challenge and possible solutions in LNG procurement activities in the wake of less availability of nuclear power capacity,” The Institute of Energy Economics, Japan (IIEJ). 2012.

- Neumann, Anne, Sophia Rüster, and Christian von Hirschhausen**, “Long-term contracts in the natural gas industry: Literature survey and data on 426 contracts (1965-2014),” 2015.
- Nunn, Nathan**, “Relationship-specificity, incomplete contracts, and the pattern of trade,” *The Quarterly Journal of Economics*, 2007, 122 (2), 569–600.
- OIES**, *The Outlook for Floating Storage and Regasification Units (FSRUs)*, Oxford Institute for Energy Studies, 2017.
- Olea, José Luis Montiel and Carolin Pflueger**, “A robust test for weak instruments,” *Journal of Business & Economic Statistics*, 2013, 31 (3), 358–369.
- Papke, Leslie E and Jeffrey M Wooldridge**, “Econometric methods for fractional response variables with an application to 401 (k) plan participation rates,” *Journal of applied econometrics*, 1996, 11 (6), 619–632.
- Ruester, Sophia**, “Changing contract structures in the international liquefied natural gas market: a first empirical analysis,” *Revue d’économie industrielle*, 2009, 127, 89–112.
- , “Financing LNG projects and the role of long-term sales-and-purchase agreements,” DIW Discussion Papers 1441, Deutsches Institut für Wirtschaftsforschung (DIW Berlin) 2015.
- Ryan, Nicholas**, “Contract enforcement and productive efficiency: Evidence from the bidding and renegotiation of power contracts in India,” *Econometrica*, 2020, 88 (2), 383–424.
- , “Holding Up Green Energy,” 2021. National Bureau of Economic Research working paper 29154.
- Songhurst, Brian**, “LNG Plant Cost Reduction 2014–18,” OIES Paper: NG137, Oxford Institute for Energy Studies, Oxford, United Kingdom October 2018.
- Steuer, Claudio**, “Outlook for Competitive LNG Supply,” OIES Paper NG 142, Oxford Institute for Energy Studies 2019.
- Sweeting, Andrew**, “Dynamic Product Positioning in Differentiated Product Markets: The Effect of Fees for Musical Performance Rights on the Commercial Radio Industry,” *Econometrica*, 2013, 81 (5), 1763–1803.
- Talus, Kim**, “Long-term natural gas contracts and antitrust law in the European Union and the United States,” *Journal of World Energy Law and Business*, 2011, 4 (3), 260–315.
- The Japan Fair Trade Commission**, “Survey on LNG Trades (Summary),” June 28. https://www.jftc.go.jp/en/pressreleases/yearly-2017/June/170628_files/170628-1.pdf. Accessed 1st December 2021. 2017.
- Weems, Philip**, “LNG Disputes on the Rise,” 2016. King & Spalding LLP report.
- Weintraub, Gabriel Y, C Lanier Benkard, and Benjamin Van Roy**, “Markov perfect industry dynamics with many firms,” *Econometrica*, 2008, 76 (6), 1375–1411.
- Williamson, Oliver E.**, “Markets and hierarchies,” *New York*, 1975, 2630.
- Williamson, Oliver E**, “Credible commitments: Using hostages to support exchange,” *American Economic Review*, 1983, 73 (4), 519–540.
- Zahur, Nahim Bin**, “Market Power and Spatial Price Discrimination in the Liquefied Natural Gas Industry,” 2023. Working Paper.

Online Appendix


Appendix A provides additional details about the industry, data and descriptive evidence. Appendix B discusses additional aspects of the estimation procedure, including the algorithm for solving for the spot market equilibrium, demand and production cost estimates, and additional estimates of contracting and investment parameters. Appendix C describes how counter-factual analysis is implemented and provides additional results relating to the counter-factuals.

A Additional Industry Details and Descriptive Evidence

A.1 Additional Industry Details

Destination clauses in long-term contracts Figure A1 illustrates how destination clauses restrict the contract buyer’s ability to resell LNG to a third party. If the long-term contract includes a destination clause, the contract buyer must take physical delivery of the LNG in a terminal in their own country, and “re-load” the cargo onto a different ship that they have to secure, which can then transport the LNG to the third party. By contrast, if no destination clause were present, the buyer could ask the seller to directly deliver the LNG cargo from the seller’s terminal to the third party (known as a “diversion”), which would both minimize shipping costs and also avoid various operational costs associated with reloading LNG, such as LNG lost through “boil-off” during the time taken to unload and reload it, as well as transaction costs associated with securing terminal capacity and an additional ship.

Figure A1: Illustrating reloads and diversions



Note: Seller A has a long-term contract with buyer 1. Each period, seller A exports an LNG cargo to buyer 1 (along Route 1) under this contract.

Reload: If a destination clause is present, Buyer 1 must take delivery of the cargo in their own home country: therefore, the only way to re-sell the cargo to buyer 2 would be to unload the cargo at their terminal, “re-load” it onto a different ship, and send the cargo to buyer 2 along Route 2.

Diversion: If no destination clause is present, then buyer 1 can ask the seller to directly send the cargo from A to buyer 2 along Route 3 (a “diversion”), which would minimize shipping costs and avoid additional operational costs of reloading LNG (such as LNG lost in boil-off during the time it takes to unload and reload LNG).

Diversion clauses in long-term contracts There are 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 (Talus, 2011). 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.⁵⁰ In its 2017 survey, Japan’s FTC found that a large share of LNG contracts featured diversion clauses (The Japan Fair Trade Commission, 2017).

Ownership structure of LNG plants 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.⁵¹ 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 because many exporting countries require that international joint ventures for LNG be either fully or majority owned by domestic firms, making full vertical integration infeasible.⁵²

Financing LNG liquefaction projects There are two main approaches for financing LNG liquefaction projects. The first is for project sponsors to finance the project themselves, either by issuing equity, or by using cash flows from existing projects or their balance sheet.

The second approach is to utilize debt finance, where the project is funded through loans from external entities (e.g., commercial banks, export credit agencies, development banks etc.). LNG projects typically use a particular type of debt finance known as *project finance*, where the project sponsors set up a project company (in the form of a special-purpose vehicle) solely to develop and operate the project (Pierru et al. (2013), Baker (2020)). It is this project company, rather than the sponsors, that receives financing and is liable for paying loans back. Lenders rely on the project’s cash flows and assets for repayment, and the sponsors themselves are generally not directly liable in the event that the project’s cash flows are insufficient; lenders only have limited recourse to the sponsors.⁵³ Because of this feature, lenders are typically unwilling to advance project finance unless the project sponsors have negotiated long-term contracts that provide a sufficient amount of guaranteed revenue to ensure the debt can be repaid (Ruester, 2015).

In practice, LNG projects often use a mix of both project financing and internal finance from the project sponsors. Between 1985 and 2020, around 40% of the total capital for major liquefaction projects was funded using project finance (Baker, 2020). *Conditional* on a project being project-financed, the debt share is on average around 67% (Natural Gas World, 2017). In recent years, the

⁵⁰Talus (2011) argues that “price-splitting” clauses often entirely remove the buyer’s incentive to resell.

⁵¹Such joint ventures are also observed in other sectors where large capital investments need to be undertaken, such as undersea fiber-optic cables (Caoui and Steck, 2023).

⁵²This is the case, for example, in Qatar and Indonesia, historically two of the world’s largest LNG exporters.

⁵³This can take the form of guarantees until project completion, commitments to complete construction and operate the plant, and prohibitions on selling interest in the project (Baker, 2020).

share of project finance has declined, as LNG projects have increasingly been financed directly by project sponsors using either equity or accrued cash flows (Baker, 2020).

Financial derivatives This paper focuses on the LNG market up until 2017, when derivatives trade played only a very limited role in the market. Since 2017, 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 was 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).

A.2 Further Details on Dataset Construction

Contract data The raw contract data is obtained from a combination of sources. The core source is a database of long-term LNG contracts maintained by Bloomberg.⁵⁴ This contains information on the identity of the seller and buyer, the liquefaction plant from which the contract is to be fulfilled, the start and end year, and the annual contract quantity. I complement this with the database of long-term LNG contracts published by Neumann et al. (2015), which covers the period from 1965 - 2014, as well as annual reports of the GIIGNL, that are available from 2004 onwards.⁵⁵ The resulting dataset consists of every long-term LNG contract signed from 2004 to 2017, as well as every long-term contract signed before 2004 that was still active in 2004 or afterwards. Finally, the dataset also includes many (though not all) long-term contracts that expired before 2004.

While the Bloomberg data does not include information on when the contract was signed, the contract signature year is available from GIIGNL for any contract signed from 2004 onwards, as well as from Neumann et al. (2015) for numerous contracts. In addition, for the majority of contracts, I hand-collected press releases, industry news reports, and snapshots from company websites, in order to verify both the contract signature date and to confirm other contract details (such as the liquefaction facility used for the contract).⁵⁶

Out of 464 long-term contracts, I was able to collect the year and month of contract signature for 425 of them. For the remaining 39 contracts (which are largely contracts that begin well before the start of the main sample period), the contract signature date is assumed to be 3 years before the start date of the contract.

Liquefaction capacity data As described in Section 2, liquefaction capacity data is obtained from the annual reports of the GIIGNL. The GIIGNL reports provide information on the charac-

⁵⁴This dataset was accessed in September 2018.

⁵⁵GIIGNL's annual reports document every new LNG contract signed in that year, as well as all past contracts that are active in that year.

⁵⁶See https://www.eneos.co.jp/english/newsrelease/jx/2011/pdf/20110510_01.pdf for an illustrative example of a press release, from JX Nippon on 10 May, 2011, announcing the signature of a 15-year long-term contract with Chevron beginning from 2015.

teristics of each plant, the project’s nameplate capacity (in millions of tons per annum), storage capacity and number of “trains”. I also observe the start-up year of each plant, the identity of the operator, and the ownership structure (i.e., the identity and shares of each of the owners). Finally, I observe whether or not the plant is a “greenfield” project: that is, whether or not the plant is built on a brand-new site rather than being added to a site that already has an existing plant. There are, however, three key variables that are not readily available from the GIIGNL reports: (i) the date of the final investment decision (FID) (ii) the cost of the investment (iii) the plant’s available capacity in any given year, distinct from its “nameplate” capacity. I now describe how I collect data on each of these variables.

Date of final investment decision (FID): For plants where the FID was reached in 2004 or after, the year and month of the FID is typically available from the annual reports of the GIIGNL. However, GIIGNL does not have information about the FID date for older plants.

Therefore, I hand-collected press releases and news articles for each LNG export terminal. Typically, the FID is a major milestone for the export project, and is therefore announced and publicized via press releases and widely publicized in news media and specialized industry outlets. From these sources, I am able to infer the year and month (and in many cases the exact date) when the FID was reached.

Investment cost: Data on the reported cost of investment is collected from numerous sources. [Songhurst \(2014\)](#) and [Songhurst \(2018\)](#) provide cost information for plants constructed between 2010 and 2018; for other plants, data on investment costs was collected from (whenever possible) press releases from project developers, as well as from [Pierru et al. \(2013\)](#), *Mechademy*, *Hydrocarbons Technology*, and old issues of *Oil & Gas Journal* and other industry publications. Overall, investment cost information is available for 53 out of 74 plants.⁵⁷

The investment cost is the cost as projected at the time when the project sponsors reach a final investment decision, which is then used to determine the total funding requirement. The realized cost of building the plant may differ significantly from the projected cost. The cost of investment also includes capitalised external finance costs ([Bartsch \(1998\)](#); [Baker \(2020\)](#)), and thus can be used to understand whether securing buyers on ex-ante long-term contracts enables sellers to obtain funding more cheaply.

Nameplate and available (or effective) capacity: LNG plants may be subject to a variety of exogenous shocks (e.g., technical failures, disruptions in gas supply to the terminal etc.) that, in any given period, may reduce the amount of liquefaction capacity available. As such, the nameplate or built capacity of a plant may not accurately reflect the true level of capacity that is available to the exporter in a given time period.

Therefore, I construct a time-varying measure of each plant’s capacity, “effective capacity”,

⁵⁷I collected investment cost information for all but 3 of the plants where an FID was reached between 1995 and 2017. For older plants with FID before 1995, reliable data on investment cost was difficult to obtain.

which better captures the capacity actually available to the plant in each year. For each LNG plant, I collect information on how many months the plant was operational for each year from 2004 - 2017, and whether the plant's production was subject to any disruption or shock. This is collected from a variety of industry sources, including a database of LNG terminal maintenance incidents published by Refinitiv Eikon, annual reports of the GIIGNL, past issues of *World Gas Intelligence*, and news articles and press releases.

The types of disruptions to LNG production that I consider, when calculating how many months the plant was operational for, include (i) maintenance and/or technical failures (ii) natural gas shortages (iii) disruptions due to war and sabotage (iv) ramp-up in production of new plants. The effective capacity of the plant is then calculated by multiplying its built capacity by the fraction of months in the year for which the plant is active. This means that if a plant experiences no disruptions or shutdowns, its effective capacity is equal to its built capacity. However, if a plant is shut down for 6 months in a year, its effective capacity is assumed to equal half its built capacity.

Table A1 shows summary statistics for nameplate and effective capacity, at the plant-year level. Out of 731 plant-year observations, 210 involve some form of production disruption experienced by the LNG plant: thus, disruptions are not uncommon.

Table A1: Summary stats for nameplate and effective capacity, at the plant-year level, measured in million tons per annum (mtpa)

Variable	Obs	Mean	Std. dev.	Min	Max
Nameplate capacity	731	4.87	2.12	0.4	10.8
Effective capacity	731	4.34	2.29	0	10.8

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.⁵⁸ 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.⁵⁹ I complement this dataset with spot price indices for North-east Asia, Singapore and Dubai/Kuwait/India from Thomson Reuters; 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.

A.3 Summary Statistics

Table A2 contains summary statistics on key variables used in the analysis. Panel A shows trade flows and shipping costs (defined at the exporter-importer-year level); Panel B illustrates spot prices

⁵⁸I accessed these data through the Reuters Eikon terminal.

⁵⁹Freight rates for exporter-importer pairs not covered by the Waterborne dataset are imputed based on a regression model linking the freight rate to the distance between two ports.

and total imports (defined at the importer-year level). Panel C show key statistics for the dataset of 464 long-term contracts. 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. 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. The average reported cost of investment is \$0.87 bn/mtpa, though there is substantial heterogeneity.⁶⁰

Table A2: 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
Reloads (mt)	2,096	0.01	0.05	0	0.78
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					
Quantity (mtpa)	464	1.28	1.10	0.04	5.20
Duration (years)	464	17.23	6.41	4	42
Time from signature to start (years)	464	3.61	2.17	0	12
Signature Year	464	2004	11.31	1963	2018
Panel D. Investment-level					
Capacity (mtpa)	74	6.94	5.02	0.50	28.90
Year of Final Investment Decision	74	2003	14.39	1959	2021
Time from FID to start date (years)	74	4.30	1.32	2.00	9.08
Average Investment Cost (\$bn/mtpa)	53	0.87	0.74	0.12	3.33

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. FID refers to the final investment decision.

Figure A2 presents a histogram of long-term contract durations. Figure A3 shows how the share of short-term contracts and spot transactions has evolved over time.

Seller and buyer capacity utilization Figure A4a shows the distribution of capacity utilization by LNG exporter - year.⁶¹ Most exporters consistently operate at a utilization over 90%, and utilization exceeds 70% for all observations. This is consistent with industry consensus that LNG export facilities try to operate at as high a level of utilization as is technically feasible (Gomes,

⁶⁰Note that this is lower than the typical costs of investment reported in recent years, since the sample includes both new and old plants, and older plants generally had a lower investment cost.

⁶¹Production (unlike capacity) is only observed at the export country-year level. So capacity utilization is calculated for each export country-year observation, as production divided by *effective* capacity.

Figure A2: Histogram of long-term contract durations





Figure A3: Share of short-term contracts and spot transactions in LNG trade



Note: Long-term contracts are contracts exceeding four years in duration (Figure A2). Figure A3 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.

2020).⁶²

By contrast, capacity utilization for LNG importers tends to be substantially lower, as Figure A4b illustrates. The average level of capacity utilization is 40%, and more than 95% of importer-year observations feature capacity utilization lower than 78%. During periods of very high demand, buyers may run into capacity constraints, as was the case in some European countries in 2022.⁶³ But such instances are rare, and in the data there are only 6 observations (out of 347) where import capacity utilization exceeds 95%. Thus, there is typically substantial excess regasification capacity available to LNG importers; suggesting that lack of regasification capacity is not a first-order constraint inhibiting LNG trade.

A.4 LNG shipping costs: spot trade vs. long-term contracts

For a subset of the study period (Dec 2010 to Dec 2017), I collect data on both spot and longer-term shipping costs from Clarksons Research. Figure A5 shows that there is no systematic difference between short-term and long-term shipping rates. The spot market is more volatile but the average freight rate is quite similar: on average, spot rates were only 2.7% higher than 1-year time-charter rates. Moreover, many spot trades are likely made on longer time-charters, because LNG operators who regularly trade on the spot can secure ships on such contracts. All of this suggests shipping costs are likely to be very similar between spot trade and long-term LNG trade.

⁶²The nameplate capacity of an LNG plant is an engineering estimate of its maximum feasible throughput under normal conditions. If conditions are favorable (e.g., planned maintenance takes less time than scheduled), production can occasionally exceed nameplate capacity, as in the case of the Equatorial Guinea LNG terminal: see https://www.lngindustry.com/lng-shipping/14022014/eg_lng_exceeds_cargo_target_164/.

⁶³<https://www.reuters.com/business/energy/brimming-european-lng-terminals-have-limited-space-more-gas-2022-02-17/>





Figure A4: Histograms of capacity utilization for LNG exporters and importers

Figure A5: LNG Freight Rates: spot vs. 1-year timecharter rates, 2010 - 17



Note: Freight rates are for 160K CBM TFDE carriers. Source: Clarksons.

A.5 Availability of alternative trading partners

This section provides additional evidence that sellers signing long-term contracts are limited in the set of buyers they can feasibly contract with. As Table A3 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 A3 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. 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, as Table A3 also shows. The relative bargaining power of the seller with respect to the buyer is an empirical question that is addressed when estimating the model.


Table A3: 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. 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.

A.6 Ex-ante contracting: additional descriptive evidence

Figure A6 illustrates heterogeneity in the share of capacity committed under ex-ante contracts: sellers typically contract a large share of capacity before investment, but some proceed with little or no ex-ante contracting. Table A4 shows that the share of capacity under ex-ante contracts for new plants has remained stable over time. Between 1998 and 2002, the share was 58.2%; between 2013 and 2017 (towards the end of the sample period), the share was 63%.

Figure A6: Distribution of share of capacity committed under ex-ante contracts**Table A4:** Average share of plant capacity committed under ex-ante contracts, for investments taking place at different points in time

Year of FID	1998 - 2002	2003 - 2007	2008 - 2012	2013 - 2017
No of new investments	10	15	13	13
Avg. share of new capacity pre-contracted	58.2%	67.0%	60.5%	63.0%

Note: This table reports the average share of plant capacity committed under ex-ante long-term contracts in the first 20 years, for LNG investments reaching final investment decisions (FID) during different time periods. The first row shows the number of new investment projects in each FID window. The second row shows the average share of plant capacity that was already committed through ex-ante contracts at the time of FID.

Relationship between geography and ex-ante contracting: further evidence 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, in that regression, is defined as the distance from the agent to its alternative trading partners (using the 1st quartile 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.⁶⁴

$$\text{Relative distance of agent } i \text{ negotiating with } j = \frac{\text{Distance from } i \text{ to alternative trading partners}}{\text{Distance between } i \text{ and } j}$$

Table A5: Contract quantity regressions: alternative measure of relative distance

Dependent variable	Quantity Share			
	(1)		(2)	
	Mean Distance		Median Distance	
Distance from Alt. Trading Partners	Estimate	S.E.	Estimate	S.E.
Ex-ante*Relative distance, seller	0.115***	(0.0417)	0.0614**	(0.0252)
Ex-ante*Relative distance, buyer	-0.111***	(0.0384)	-0.0611**	(0.0242)
Relative distance, seller	-0.015	(0.021)	-0.0088	(0.015)
Relative distance, buyer	0.016	(0.018)	0.012	(0.014)
Ex-ante contract	0.051	(0.032)	0.056*	(0.031)
Other Controls	Distance, Capacity, Extension, Trend, Greenfield			
N	337		337	
R ²	0.27		0.27	

Note: Each observation is a long-term contract. The dependent variable is the contract quantity expressed as a share of plant capacity. The relative distance of an agent is defined as their distance from alternative trading partners divided by the distance from their current trading partner. The numerator, the distance to alternative trading partners, is defined as the mean distance between that agent and all potential trading partners (in (1)), and the median distance between that agent and all potential trading partners (in (2)). Other controls include the distance between the seller and buyer negotiating the contract, the logarithm of plant capacity, an indicator for greenfield plants, an indicator for contract extensions, and a time trend. Statistical significance at the 10%, 5%, and 1% levels are denoted with *, **, and ***, respectively.

Table A5 repeats the contract quantity regression shown in Table 2, but using two alternative measures of relative distance that differ in how I measure the agent's proximity to potential trading

⁶⁴An example may be helpful: 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., there are many other buyers in Europe).

partners. Instead of using the 1st quartile of the distance (as in the original regression in Table 2), I measure how far an agent is from alternative trading partners by the mean of the distance from all potential trading partners, as well as the median. The results are similar.

Relationship between geography and ex-ante contracting at the plant level I also investigate the relationship between the strength of outside options and ex-ante contracting at the plant level, instead of at the contract level as in the preceding regressions. For each export project, I compute the share of its total capacity committed under ex-ante contracts in the first 20 years, and regress it on the distance from the seller to alternative buyers (which, as before, is defined as the 1st quartile of the distance from the seller to all buyers). As Table A6 shows, the coefficient on distance to alternative buyers is positive and significant and is robust to controlling for the project capacity, a time trend as well as an indicator for greenfield projects.⁶⁵ The magnitude of the coefficient suggests that increasing the distance between an export project and alternative buyers by 1000 nautical miles will increase the share of capacity contracted ex-ante by 7-11 percentage points.

Table A6: Regression of share of total capacity committed under ex-ante contracts on project features

	(1)	(2)	(3)	(4)
	Dependent variable: Share of capacity committed under ex-ante contracts			
Distance from alternative buyers	0.069* (0.039)	0.084** (0.040)	0.11** (0.048)	0.11** (0.047)
log(Capacity)		-0.12* (0.071)	-0.14* (0.073)	-0.12 (0.072)
Time Trend			-0.0079 (0.0084)	-0.0098 (0.0083)
Greenfield				0.18* (0.099)
N	58	58	58	58
R ²	0.052	0.099	0.11	0.16

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 alternative buyers is measured in 1,000 nautical miles. Statistical significance at the 10%, 5%, and 1% levels is denoted with *, **, and ***, respectively.

A.7 Additional descriptive evidence of contract rigidity

Table 4 of Section 3 shows, using a fractional logit regression, that sellers with a high share of contracted capacity are less responsive to price differentials. The dependent variable in that regression, s_{ijt} , is a fractional variable bounded between 0 and 1. Column (1) of Table A7 show that this con-

⁶⁵Table A6 suggests that sellers rely more on ex-ante contracting for greenfield plants, though the estimate is rather noisy.

clusion is unchanged if I were to ignore the fractional nature of the dependent variable and directly regress s_{ijt} on the net price p_{ijt} (instrumenting for p_{ijt}). Columns (2) and (3) show that I obtain similar results under different transformations of the dependent variable (e.g., inverse hyperbolic sine (IHS) transformation, or adding 1 to the shares and taking logarithms).

Table A7: Responsiveness of allocations to price differentials: further evidence

	(1)	(2)	(3)	(4)
Dependent Variable	s_{ijt}	IHS-transformed s_{ijt}	$\log(s_{ijt} + 1)$	Rank of s_{ijt}
Methodology	2SLS	2SLS	2SLS	OLS
Net price, p_{ijt}	0.045*** (0.0078)	0.043*** (0.0075)	0.038*** (0.0065)	
Net price* Share Contracted	-0.047*** (0.0097)	-0.046*** (0.0092)	-0.041*** (0.0079)	
Rank of net price				0.29*** (0.042)
Rank of net price * Share Contracted				-0.25*** (0.072)
Fixed Effects	Seller-Year and Seller-Buyer FE			
N	3116	3116	3116	3246
1st-stage F-stat (Cragg-Donald)	16.9	16.9	16.9	
1st-stage F-stat (Kleibergen-Paap)	11.9	11.9	11.9	

Note: Each observation is an exporting country (seller) - importing country (buyer) - year. The net price seller i receives from selling to buyer j , p_{ijt} is the difference between the spot price paid by buyer j , p_{jt} , and the the per-unit cost of shipping LNG from i to j , c_{ijt}^d is: $p_{ijt} = p_{jt} - c_{ijt}^d$. "Share Contracted" refers to the share of the seller's capacity that committed under long-term contracts.

Columns (1) - (3) explore different ways to define the dependent variable. In Column (1), the dependent variable is s_{ijt} , the share of seller i 's total LNG production that is sold to buyer j ; and the parameters are estimated using 2SLS. In Column (2), the dependent variable is the inverse hyperbolic sine (IHS) of the share s_{ijt} . In Column (3), the dependent variable is $\log(s_{ijt} + 1)$. In Columns (1) - (3), to account for endogeneity of p_{ijt} , I instrument for p_{ijt} by shipping costs (c_{ijt}^d) and demand shifters that affect country j 's demand for LNG in period t (specifically, j 's electricity consumption from fossil fuels, coal price and the minimum temperature in country j). In Columns (1) - (3), the dependent variable, s_{ijt} , is the share of seller i 's total LNG production that is sold to buyer j .

In Column (4), the dependent variable is the rank of s_{ijt} within that seller-year: the rank is highest for the buyer j to whom seller i allocates the largest share of their output. The key regressor is the rank of net price p_{ijt} within that seller-year: the rank is highest for the buyer j from whom the seller can obtain the highest price net of shipping costs.


Building on these results, Column (4) investigates whether a seller's allocation of output across buyers is sorted by net price and how this sorting is influenced by the share of the seller's capacity committed under long-term contracts. To do this, I first calculate, for each seller i in period t , the rank of each buyer based on the share of output the seller allocates to that buyer (s_{ijt}), as well as the rank of each buyer based on the net price (p_{ijt}). Then I regress the rank of s_{ijt} on the rank of net price (p_{ijt}), allowing the coefficient to vary with the share of the seller's capacity that is contracted. As with all the other specifications, I continue to include seller-year and seller-buyer fixed effects. I

find a positive coefficient on the rank of price, indicating that buyers from whom the seller receives a higher net price are more likely to be allocated a larger share of the seller's output. However, the coefficient on the interaction term between rank of price and the share contracted is negative, suggesting that contractual rigidities dampen this sorting effect.

A.8 Evolution of spot prices and trade flows in LNG industry

Here I describe the evolution of spot prices and trade flows in the industry, and argue that these patterns are suggestive of contract rigidities. Figure A7 plots spot prices in various LNG importing regions, showing there are 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 A7: 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 frictions, faced with these divergent prices, capacity-constrained sellers might be expected to concentrate their sales on the destinations with the highest price net of transportation costs. This would mean that during these periods of high Asia-Europe spot price differentials, LNG exports (if competitively allocated) 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 during this period, Europe continued to import large amounts of contracted LNG,

as shown by Appendix Figure A8b, 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 have impeded the market's response to the demand shock.⁶⁶

Figure A8: LNG imported by different regions



Note: Long-term contracts are defined as contracts exceeding four years in length. Short-run and spot LNG imports use spot transactions or contracts that are less than or equal to four years in length. Source: GIIGNL.

B Estimation Details

B.1 Algorithm for solving spot market equilibrium in baseline model

In the baseline spot market model, sellers face binding capacity constraints and are Cournot competitors. Under the assumption of linear production cost, the profit maximization problem of any seller i on the spot market can be written as:

$$\{S_{ijt}\}_{j=1}^J = \underset{\{\hat{S}_{ijt}\}_{j=1}^J}{\operatorname{argmax}} \left[\sum_{j=1}^J p_{jt}^*(\hat{S}_{ijt}, S_{-ijt}) \hat{S}_{ijt} - \sum_j c q_{ijt} - \sum_j c_{ijt}^d q_{ijt} \right]$$

$$\text{s.t. } q_{it} \leq K_{it}$$

Part of the seller's production is used to fulfill contracts. Let $\kappa_{it} = K_{it} - \sum_j q_{ijt}^c$ denote the seller's total *uncommitted capacity* in period t . Then the capacity constraint can be equivalently written as:

$$\sum_j S_{ijt} \leq \kappa_{it} \quad (\text{B1})$$

The equilibrium in period t is characterized by a set of spot quantity choices by each seller i ,

⁶⁶That is not the whole story, however, as some sellers (e.g., Qatar) with a higher net price from selling to Asia continued to sell *spot cargoes* to Europe. This suggests LNG spot sellers exercise market power, which is further explored by [Zahur \(2023\)](#).

$\{S_{ijt}\}_j^J$, such that each seller's capacity constraint (B1) and first-order condition (B2) is satisfied:

$$\underbrace{p_{jt}^* + S_{ijt} \frac{\partial p_{jt}^*(S_{ijt}, S_{-ijt})}{\partial S_{ijt}}}_{\text{Marginal revenue of selling to market } j} - \underbrace{\left(\frac{\partial C(q_{it}, K_{it})}{\partial S_{ijt}} + \lambda_{it} + c_{ijt}^d \right)}_{\text{Marginal opportunity cost of selling to market } j} \leq 0 \quad (\text{B2})$$

with equality if $S_{ijt} > 0$. Note that the Lagrange multiplier λ_{it} equals 0 if $\sum_j S_{ijt} < \kappa_{it}$.

To solve for this multi-regional spot market equilibrium, I rely on the results from [Alsabab et al. \(2021\)](#), and adapt their numerical algorithms. The key result of [Alsabab et al. \(2021\)](#) is that there is a one-to-one mapping between the (unique) Cournot equilibrium, where sellers choose allocations (S_{ijt}) , and the Nash equilibrium of a reduced-form game, where sellers instead choose their “marginal profits” (the Lagrange multiplier λ_{it}) to maximize their own profits. Thus, solving for the Nash equilibrium of this reduced-form game allows us to recover the Cournot equilibrium.

Based on this result, [Alsabab et al. \(2021\)](#) develop two algorithms that can together be used to solve for the unique Cournot equilibrium in this multi-region setting (both of which I describe below). In practice, to compute the spot market equilibrium, I first run Algorithm 2, which solves for the equilibrium marginal profits in the reduced-form representation of the game; and then I run Algorithm 1 to solve for the corresponding spot market allocation.

Algorithm 1: recover allocations from marginal profits: For any vector of marginal profits $\lambda_t = \{\lambda_{it}\}_{i=1}^N$, this algorithm can be used to solve for the corresponding spot market allocation, $\{S_{ijt}\}_{i,j}$, in each period t (and is guaranteed to find the allocation in a finite number of steps).

Before describing the formal algorithm, it will be helpful to define a few objects. First, for each buyer j , I rank sellers based on their competitiveness, which is the inverse of their *opportunity cost* of selling an additional unit to that buyer, $\phi_{jt}(i)$. The opportunity cost is the sum of marginal production costs, shipping costs, and the marginal profit: $\phi_{jt}(i) = c_i + c_{ijt}^d + \lambda_{it}$. Let σ_{jt} be an ordering of sellers such that $\phi_{jt}(i)$ is weakly increasing in $\sigma_{jt}(i)$. Thus, sellers with a lower ranking have a lower opportunity cost ϕ and are more competitive when selling to j . Second, let I_{jt} denote the set of all sellers who sell strictly positive quantities to buyer j .

I now show that if both I_{jt} and marginal profits λ_{it} are known for each seller i and each buyer j , then I can recover the spot market allocations $\{S_{ijt}\}_{i,j}$. Let $S_{jt}^{tot} = \sum_i S_{ijt}$ denote buyer j 's total spot imports in period t . Now observe that for a seller i who belongs to the set I_{jt} , the FOC (equation (B2)) must hold with equality:

$$p_{jt}(S_{jt}^{tot}) + p'_{jt}(S_{ijt})S_{ijt} - c - c_{ijt}^d - \lambda_{it} = 0 \quad (\text{B3})$$

Add up (B3) across all k firms that belong to I_{jt} :

$$kp_{jt}(S_{jt}^{tot}) + p'_{jt}(S_{ijt})S_{jt}^{tot} - \sum_{i \in I_{jt}} (c + c_{ijt}^d + \lambda_{it}) = 0$$

$$kp_{jt}(S_{jt}^{tot}) + p'_{jt}(S_{jt}^{tot})S_{jt}^{tot} = \sum_{i \in I_{jt}} (c + c_{ijt}^d + \lambda_{it}) \quad (\text{B4})$$

Define: $G_{jt,k}(S_{jt}^{tot}) = kp_{jt}(S_{jt}^{tot}) + p'_{jt}(S_{jt}^{tot})S_{jt}^{tot}$. Since this function is strictly monotonic in S_{jt}^{tot} , it can be inverted to solve for S_{jt}^{tot} :

$$S_{jt}^{tot} = G_{jt,k}^{-1} \left(\sum_{i \in I_{jt}} (c + c_{ijt}^d + \lambda_{it}) \right) \quad (\text{B5})$$

Once S_{jt}^{tot} has been solved for, S_{jt} can be recovered using the following equation that can be derived by re-arranging equation (B2):

$$S_{ijt} = \frac{p_{jt}(S_{jt}^{tot}) - c - c_{ijt}^d - \lambda_{it}}{-p'_j(S_{ijt})} \text{ if } i \in I_{jt} \quad (\text{B6})$$

Now we are in a position to describe Algorithm 1:

Algorithm 1: Initialize $I_{jt}^{(0)} = \emptyset$, and $S_{jt}^{tot,(0)} = 0$, for each market j . Then at iteration k , $k \geq 1$:

1. If $k = N + 1$ (we have exhausted all sellers) or if $\phi_{jt}(k) \geq p_{jt}(S_{jt}^{tot,(k-1)})$ (next most competitive seller doesn't want to sell), then $I^{(k)} = I^{(k-1)}$ and the algorithm is terminated. $S_{jt}^{tot,(k-1)}$ is now already known, and we can use equation (B6) to solve for each firm's sales S_{ijt} .
2. Otherwise, set $I^{(k)} = I^{(k-1)} \cup k$ (that is, we add the next most competitive firm). Then, use equation (B5) to solve for total sales to j , $S_{jt}^{tot,(k)}$, and proceed to the next step.

Algorithm 2: solve for equilibrium of reduced-form game where firms choose marginal profits:

The fixed point algorithm described below can be used to solve for the equilibrium of the reduced-form game where sellers compete by choosing marginal profits.⁶⁷

Best response function: Let $\lambda_{-it} = (\lambda_{1t}, \dots, \lambda_{i-1,t}, \lambda_{i+1,t}, \dots, \lambda_{Nt})$ denote a vector containing marginal profits for all sellers except i . Let $q_{it}^{tot}(\lambda_{it}, \lambda_{-it})$ denote seller i 's total spot sales as a function of their own marginal profit and rival marginal profits: note that this can be derived by applying Algorithm 1. $q_{it}^{tot}(\lambda_{it}; \lambda_{-it})$ is monotonically decreasing in λ_{it} (the higher the opportunity cost, the lower the seller's sales), and therefore this function can be inverted. The best response function for firm i (when choosing λ_{it} in response to λ_{-it}) can then be written as:

$$\Phi_{it}(\lambda_{-it}) = \begin{cases} 0 & \text{if } 0 \leq q_{it}^{tot}(0; \lambda_{-it}) < \kappa_{it} \\ (q_{it}^{tot})^{-1}(\kappa_{it}; \lambda_{-it}) & \text{if } \kappa_{it} \leq q_{it}^{tot}(0; \lambda_{-it}) \end{cases} \quad (\text{B7})$$

⁶⁷ Alsabah et al. (2021) prove that if sellers iteratively choose their best responses to the marginal profits chosen by other sellers, marginal profits will converge to the equilibrium of the reduced-form game. Therefore, a fixed point algorithm can be used to numerically compute the equilibrium.

The intuition is as follows: given the marginal profits λ_{-it} chosen by its rivals, if firm i were not capacity constrained, they would best respond by producing a total quantity such that their marginal profit was zero (since they could not raise their profits by increasing production any further): this quantity equals $q_{it}^{tot}(0; \lambda_{-it})$. If this quantity is lower than firm i 's uncommitted capacity κ_{it} , then firm i 's best response is to choose zero marginal profit. But if firm i does not have enough capacity to produce $q_{it}^{tot}(0; \lambda_{-it})$, then its optimal choice of marginal profit will be such that its production exactly equals its uncommitted capacity, or $q_{it}^{tot}(\lambda_{it}; \lambda_{-it}) = \kappa_{it}$; which can be inverted to yield its best response, $\lambda_{it} = (q_{it}^{tot})^{-1}(\kappa_{it}; \lambda_{-it})$.

Fixed point algorithm:

1. Start with an initial guess of the vector of marginal profits λ_t : call this $\lambda_t^{(0)}$.
2. Update the guesses of the marginal profits. At each iteration l , for every seller i , update λ_{it}^l using the best response operator equation (B7), taking as given the marginal profits chosen by all other sellers in the $(l-1)$ th iteration, or $\lambda_t^{(l-1)}$.
3. Stop iterating once $\|\lambda_{it}^{(l)} - \lambda_{it}^{(l-1)}\| < tol$ for each i , where tol is a pre-assigned tolerance level.

B.2 Demand Estimates: Additional Details and Results

This section provides further details on demand estimation. Table B8 shows the baseline demand estimates. The coefficient on electricity consumption from fossil fuels (which is the elasticity of LNG demand with respect to residual electricity demand) is 1.11. The elasticity of LNG demand with respect to the oil price is 0.80 for the median observation. Finally, the coefficient on the minimum temperature implies that a one-standard deviation decrease in the minimum temperature would cause LNG demand to increase by 9%. Thus, LNG demand is quite responsive both to changes in temperature and shifts in electricity demand, as well as to changes in the price of oil (a substitute for natural gas).

The demand system in equation (11) does not explicitly take into account capacity constraints faced by the buyer. These capacity constraints, however, are very rarely close to binding: import capacity utilization is typically quite low (40% on average), and is less than 90% for more than 94% of the observations.⁶⁸ Dropping the handful of observations where capacity utilization is more than 90% leads to almost identical demand estimates.

⁶⁸Moreover, under certain circumstances, countries can (if they need to) import LNG in excess of their capacity, as was the case for Taiwan from 2012 to 2017.

Table B8: Demand Curve Estimates

	1st-stage		2SLS	
	Estimate	S.E.	Estimate	S.E.
Spot Price			-0.106	(0.055)
log(Elec. Cons., fossil)	0.63	(0.51)	1.115	(0.201)
Min. temp	-0.23	(0.098)	-0.092	(0.027)
Oil Price	0.089	(0.0041)	0.018	(0.0056)
<i>Excluded IVs:</i>				
Min. temp. else	-2.09	(0.74)		
Elec cons. else, fossil	0.0012	(0.00027)		
Elec gen. else, baseload	-0.0018	(0.00038)		
Olea-Pflueger F-stat	10.1		Median elasticity	-0.83
Kleibergen-Paap F-stat	9.83		Mean elasticity	-0.95

Note: Each observation is an importer-year-quarter pair. There are 815 observations. 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 is the logarithm of total LNG imports in country j in period t . The minimum temperature is standardized so that it has zero mean and a variance of one. All regressions include importing country and quarter fixed effects. The instruments for prices are the average of the normalized minimum temperature in period t for all importing countries excluding country j ("Min temp else"), total electricity consumption from fossil fuels in period t excluding country j 's own consumption ("Elec cons. else, fossil"), and total electricity generation from baseload sources (nuclear, renewables, hydro) in period t excluding country j ("Elec gen. else, baseload"). Effective F-statistics for excluded IVs following [Olea and Pflueger \(2013\)](#) are reported, as well as the Kleibergen-Paap F-statistic. Robust standard errors in parentheses.

B.3 Cost Estimates


Baseline cost estimates Table [B9](#) presents cost function estimates for the baseline specification with binding capacity constraints (where marginal costs are constant and equal to c until the seller hits capacity); these estimates were discussed in [Section 5.2](#).

Table B9: Cost Parameter Estimates: Baseline Specification

Cost Parameters	Estimate	S.E.		
Marginal cost, c	5.42	(0.0068)		
Fit (R^2)				
Prices p_j	0.87		Production q_i	0.98
Spot Trade Flows S_{ij}	0.22		Regional Spot Trade Flows	0.51

Note: Marginal costs are constant and equal to c (which is estimated), up until the seller hits their binding capacity constraint K_{it} . 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.

Figure B9: Actual vs. Predicted Spot Prices



Note: The figure plots actual spot prices as well as model-predicted annual spot prices in Asia and Europe. The spot price in Asia is calculated as a quantity-weighted average of the spot price paid by each individual importing country in Asia; likewise for Europe.

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 0.87. Figure B9 shows actual versus predicted prices in Asia and Europe, illustrating that the model does well at matching year-to-year changes in prices. The R^2 for total exports is 0.98: this is partly because total exports are largely determined by long-term contracts, but the model also does well at predicting total spot sales of each seller. The R^2 for spot trade flows is on the low side at 0.22, but the model provides a good fit ($R^2 = 0.51$) for spot trade flows aggregated to the regional level (e.g., spot exports from Middle East to Northeast Asia). This is because buyers within the same region - such as Japan and South Korea - are geographically close and face nearly identical spot prices, making sellers largely indifferent between them. As a result, the model finds it difficult to predict the exact allocation of spot LNG between buyers in the same region, but captures regional patterns well.

Alternative specification of cost function Here I explore a specification of the cost function where capacity constraints are “soft” rather than binding (Besanko and Doraszelski, 2004):

$$C(q_{it}, K_{it}) = \delta \frac{1}{1 + \nu} \left(\frac{q_{it}}{K_{it}} \right)^\nu q_{it} \quad (\text{B8})$$

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. Unlike the baseline model, though, capacity constraints here are soft: the firm can, if it wishes, produce q_{it} above K_{it} .

I found it is difficult to identify ν and δ separately from one another.⁶⁹ As such I consider two

⁶⁹This is because high levels of capacity utilization can be rationalized both by high ν and high δ .

different calibrated values of ν : $\nu = 1$, implying marginal cost is linear in capacity utilization; and $\nu = 2$, where marginal cost is quadratic in capacity utilization. Table B10 shows the estimates. In column (1), where marginal cost is linear in capacity utilization, δ is estimated to be 14.6; while in column (2), when marginal cost is quadratic in capacity utilization, δ is estimated to be 14.8. Thus, both models suggest the marginal cost of production at 100% capacity utilization is around \$15/MMBtu, which is significantly higher than the average spot price of \$8.7/MMBtu.⁷⁰ However, the model in column (2) with $\nu = 2$ implies the marginal cost changes more sharply with capacity utilization, so that it would equal only \$3.7/MMBtu when a firm operates at 50% capacity utilization (versus \$7.3/MMBtu for the model in column (1) with $\nu = 1$).

Both these models with a soft capacity constraint provide a reasonably good fit of the data, though the model with $\nu = 2$ delivers an overall better fit. Comparing Table B10 with the baseline cost estimates in Table B9, we see that the baseline model does slightly worse at fitting regional spot trade flows, but delivers a better fit of both LNG prices and total production compared to the two models with soft capacity constraints.

Table B10: Cost Parameter Estimates: Soft Capacity Constraints

	(1) MC linear in capacity utilization	(2) MC quadratic in capacity utilization
δ	14.56 (0.049)	14.82 (0.054)
ν (calibrated)	1	2
Fit (R^2)		
Prices (p_j)	0.67	0.78
Total Exports by Exporter, q_i	0.94	0.96
Spot Trade Flows, S_{ij}	0.29	0.36
Regional Spot Trade Flows, $S_{ij\bar{}}$	0.42	0.56

Note: Each observation is an exporter-importer-year pair ($N = 3245$). The cost function follows equation (B8). In Column (1), the parameter ν is calibrated to 1 (i.e., MC is linear in capacity utilization). In Column (2), ν is calibrated to 2 (i.e., MC is quadratic in capacity utilization). 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.

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

This section presents additional estimates of the parameters characterizing contracting and investment behavior. First, I investigate the role of costly financing in driving the results. In Specification (2) of Table B11, I present parameter estimates when γ_3 is set to 0, so that ex-ante contracts do not affect financing costs; for comparison, the baseline results are shown again in Specification

⁷⁰In this model of soft capacity constraints, δ can be interpreted as the marginal cost of production for a firm operating at 100% capacity utilization, since the marginal cost equals $\delta \left(\frac{q_{it}}{K_{it}}\right)^\nu = \delta$ when $q_{it} = K_{it}$, regardless of the value of ν .

(1). The parameter governing the first-stage contract premium, θ_1 , increases from \$0.26/MMBtu to \$0.64/MMBtu, while the other parameters remain very robust. This suggests that when the financing channel is shut down, the model attributes a greater share of the ex-ante contracting to a higher first-stage contract premium.

Table B11: Contracting and investment parameter estimates: with and without costly financing

	(1)		(2)	
	Costly Financing (Baseline)		No Financing Cost	
	Estimate	S.E.	Estimate	S.E.
γ_1	31.95	(7.77)	35.50	(7.97)
$\gamma_1 \times \mathbb{1}\{\text{Greenfield}\}$	10.25	(7.77)	9.85	(7.91)
γ_2	-0.16	(0.13)	-0.12	(0.18)
γ_3	3.50	(4.66)	0	
τ	0.63	(0.09)	0.63	(0.09)
κ_1	0.0062	(0.0055)	0.0031	(0.0053)
θ_1	0.26	(0.40)	0.64	(0.40)
κ_3	-0.0012	(0.0022)	-0.0012	(0.0022)
θ_3	0.96	(0.06)	0.96	(0.06)

Note: Specification (1) is the baseline specification, where I allow ex-ante contracts to affect the cost of finance (through the parameter γ_3 . In specification (2), I set γ_3 to equal 0. All parameters estimated using non-linear least squares. Standard errors are heteroskedasticity-robust.

In Table B12, I report the results from specifications where the Nash bargaining weight τ is allowed to vary across seller and buyer groups.⁷¹ 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. While the major NOCs have a higher bargaining weight, the difference relative to other sellers is relatively small and statistically insignificant; 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; but I find that τ is similar for both groups of buyers. 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 sizable and are difficult to estimate with precision. As such, I use the baseline specification with a single bargaining weight τ for the main analysis.

Finally, 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). In Column (1)

⁷¹Since 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 B12: Contracting and investment parameter estimates: determinants of bargaining power

	(1)		(2)		(3)	
	Estimate	S.E.	Estimate	S.E.	Estimate	S.E.
γ_1	31.98	(7.77)	31.96	(7.79)	31.92	(7.81)
$\gamma_1 \times \mathbb{1}\{\text{Greenfield}\}$	10.25	(7.77)	10.25	(7.77)	10.25	(7.78)
γ_2	-0.15	(0.13)	-0.16	(0.13)	-0.15	(0.13)
γ_3	3.44	(4.67)	3.49	(4.80)	3.50	(4.79)
κ_1	0.0065	(0.0054)	0.0062	(0.0056)	0.0061	(0.0055)
θ_1	0.24	(0.40)	0.26	(0.40)	0.25	(0.40)
κ_3	-0.0012	(0.0022)	-0.0012	(0.0022)	-0.0012	(0.0022)
θ_3	0.96	(0.06)	0.96	(0.06)	0.96	(0.06)
Bargaining weight, τ :						
$\tau(\text{Other exporters})$	0.54	(0.28)				
$\tau(\text{Major NOCs})$	0.68	(0.26)				
$\tau(\text{Other importers})$			0.62	(0.79)	0.61	(1.36)
$\tau(\text{NE-Asian importers})$			0.63	(0.37)		
$\tau(\text{NE-Asian importers*Other exporters})$					0.53	(0.48)
$\tau(\text{NE-Asian importers*Major NOCs})$					0.67	(0.44)

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). “NE-Asian importers” includes Japan, China, South Korea and Taiwan. All parameters estimated using non-linear least squares. Standard errors are heteroskedasticity-robust.

of Table B13, I allow the contract premium to differ for buyers located in the Asia-Pacific region, who, unlike buyers in Europe and North America, tend to have less access to pipeline gas, and may in principle have different preferences for long-term contracts; however, I find no significant difference in their contract premium compared to other buyers. Column (2) 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). I find little evidence that the rule of law affects either ex-post or ex-ante contracting decisions. One possible explanation is that there are large reputational costs of breaching long-term LNG contracts (as discussed in Section 2), diminishing the additional benefits of contracting with countries that have stronger legal enforcement. Finally, column (3) includes an indicator for whether the buyer and seller contracted in the past, finding that it makes little difference to the contract premium. In general, these estimates suggest that heterogeneity in the contract premium across buyers is unlikely to be of first-order importance.

Table B13: Contracting and investment parameter estimates: determinants of contract premium

	"(1)"		"(2)"		"(3)"	
	Estimate	S.E.	Estimate	S.E.	Estimate	S.E.
γ_1	32.37	(7.88)	32.51	(7.80)	32.49	(7.80)
$\gamma_1 \times \mathbb{1}\{\text{Greenfield}\}$	10.21	(7.75)	10.19	(7.74)	10.20	(7.77)
γ_2	0.63	(0.09)	0.69	(0.12)	0.69	(0.12)
γ_3	3.23	(4.71)	3.33	(4.81)	3.34	(4.82)
τ						
Contract premium (ex-ante)						
κ_1	0.0074	(0.0058)	0.0049	(0.0061)	0.0052	(0.0065)
θ_1	0.55	(0.57)	1.07	(0.91)	1.05	(0.92)
θ_1 : Importer, Pacific	-0.44	(0.40)	-0.36	(0.43)	-0.39	(0.51)
θ_1 : rule of law, importer			-0.34	(0.29)	-0.35	(0.31)
θ_1 : rule of law, exporter			-0.22	(0.14)	-0.23	(0.14)
θ_1 : contracted in the past					0.07	(0.40)
Contract premium (ex-post)						
κ_3	-0.0014	(0.0022)	-0.0034	(0.002)	-0.0032	(0.0022)
θ_3	0.84	(0.17)	1.09	(0.41)	1.07	(0.52)
θ_3 : Importer, Pacific	0.17	(0.16)	0.13	(0.16)	0.13	(0.17)
θ_3 : rule of law, importer			-0.20	(0.21)	-0.21	(0.23)
θ_3 : rule of law, exporter			0.07	(0.05)	0.07	(0.05)
θ_3 : contracted in the past					0.03	(0.35)

Note: "Importer, Pacific" is an indicator for buyers in the Asia-Pacific region (located in Asia, Middle East or South America). All parameters estimated using non-linear least squares. Standard errors are heteroskedasticity-robust.

C Counter-factual Details

C.1 Algorithms for solving spot market equilibrium without resale restrictions

I describe here how I model the spot market when there are no resale restrictions, as well as the set of numerical algorithms I use to solve for the spot market equilibrium (building on the discussion in Section 6.2).

Recall that the key assumption I make is that absent resale restrictions, sellers are unable to engage in price discrimination, since any attempt to do so can be arbitrated away by buyers with long-term contracts (who now have the option of reselling LNG they import under long-term contracts). As such, in the resulting Cournot equilibrium, each seller i (simultaneously) only chooses their total spot production S_{it} (i.e. their total production above their contractual commitment); the *spatial allocation* of LNG is perfectly competitive.

C.1.1 Characterizing the spot market equilibrium with no resale restrictions

Characterizing seller choices of S_{it} , or their total spot output In a Cournot equilibrium with no resale restrictions, each seller i takes as given rival optimal strategies $\{S_{-it}\}$ and chooses S_{it} that maximizes their total profits, subject to the constraint that their *total* production (sum of contracted and spot output) cannot exceed capacity:

$$\begin{aligned} \max_{S_{it}} & \left[p_{it}(S_{it}, S_{-it})S_{it} - C(q_{it}, K_{it}) \right] \\ \text{s.t. } & q_{it} = S_{it} + \sum_j q_{ijt}^c \leq K_{it} \end{aligned}$$

Here p_{it} , the net price received by seller i on its spot production, equals:

$$p_{it}(S_{it}, S_{-it}) = \max_j (p_{jt} - c_{ijt}^d)$$

The Lagrangian is given below, with λ_{it} denoting the Lagrange multiplier on the capacity constraint:

$$\max_{S_{it}} \left[p_{it}(S_{it}, S_{-it})S_{it} - C(q_{it}, K_{it}) + \lambda_{it}(K_{it} - S_{it} - \sum_j q_{ijt}^c) \right]$$

The first-order condition satisfied by the optimal spot quantity S_{it} is:

$$p_{it} + S_{it} p'_{it}(S_{it}, S_{-it}) - \frac{\partial C(q_{it}, K_{it})}{\partial S_{it}} - \lambda_{it} \leq 0 \quad (\text{C9})$$

with equality iff $S_{it} > 0$, and $\lambda_{it} = 0$ if $q_{it} < K_{it}$.

Characterizing equilibrium spatial allocations, S_{ijt} Let $q_{it} = S_{it} + \sum_j q_{ijt}^c$ denote seller i 's total production in any given period. Let $Q_t = \{q_{it}\}_{i=1}^N$ denote a vector recording the total production of each seller in period t . Given any Q_t , let $A_t(Q_t) = \{q_{ijt}^*\}_{i=1, j=1}^{N, J}$ denote the corresponding *competitive* spatial allocation of LNG – defined as the perfectly competitive allocation subject to the constraint that each seller supplies no more than their chosen production level, q_{it} .

Following [Samuelson \(1952\)](#), this allocation can be determined by solving a social planner's problem, where the planner maximizes social surplus, subject to the constraint that each seller's production does not exceed their chosen level. This planner's problem can be written as:

$$\begin{aligned} \max_{\{q_{ijt}\}_{i=1, j=1}^{N, J}} & \underbrace{\sum_j CS_{jt}(Q_{jt})}_{\text{Buyer surplus}} - \underbrace{\sum_i C(q_{it}, K_{it})}_{\text{Seller production cost}} - \underbrace{\sum_{ij} c_{ijt}^d q_{ijt}}_{\text{Shipping cost}} \\ \text{s.t. } & q_{ijt} \geq 0 \end{aligned} \quad (\text{C10})$$

$$\text{and } \sum_j q_{ijt} \leq S_{it} + \sum_j q_{ijt}^c$$

where Q_{jt} is buyer j 's total imports in period t , and $CS_{jt}(Q_{jt})$ is the total surplus earned by buyer j from purchasing a total quantity of Q_{jt} (equivalently, buyer j 's consumer surplus from Q_{jt} units if the price were zero, or the area under buyer j 's demand curve for a quantity Q_{jt}).

C.1.2 Numerical algorithm

I use the following numerical algorithm to solve for the spot market equilibrium described above:

1. Start with an initial guess of the vector of spot production levels chosen by each seller, $\{S_{it}^{(0)}\}_{i=1}^N$.
2. At each iteration l , update the guess of the vector of spot production levels from $\{S_{it}^{(l-1)}\}_{i=1}^N$ to $\{S_{it}^{(l)}\}_{i=1}^N$ as follows:
 - (a) Solve for the competitive spatial allocation of LNG, $\{q_{ijt}^*\}_{i=1, j=1}^{N, J}$, given that seller spot production levels are $\{S_{it}^{(l-1)}\}_{i=1}^N$, by numerically solving the planner's problem (equation (C10)).⁷²
 - (b) Given the above competitive allocation, calculate, for each seller i , p_{it} (net price received by seller i) as well as p'_{it} (derivative of net price with respect to S_{it}).
 - (c) Invert each firm i 's first-order condition characterizing their optimal spot quantity, equation (C9), to obtain an updated guess of $S_{it}^{(l)}$ for every seller. Then our new guess is $\{S_{it}^{(l)}\}_{i=1}^N$.
3. Stop iterating once $\|S_{it}^{(l)} - S_{it}^{(l-1)}\| < tol$ for each i , where tol is a pre-assigned tolerance level.

C.2 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 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

⁷²In practice, I "linearize" the demand system by approximating each buyer's demand as a linear function of the price (with buyer-specific coefficients); in that case, the planner's problem can be expressed as a quadratic programming problem and can be quickly and efficiently solved using standard solvers.

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 have much effect on 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 $\mathbf{Y}_{-i}$ 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 (6), (9), and (10). 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 (9), is satisfied. Finally, in counter-factuals with no *ex-ante* contracting permitted, I search numerically for the investment K_i and the ex-post contract quantity $q_{ij}^{c,3}$ that satisfy the first-order conditions (6) and (9).

In the general equilibrium counter-factuals described in Section 6.2, 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 $\mathbf{Y}_{-i}^{l-1} = \{\mathbf{q}_{-i}^{c,l-1}, \mathbf{K}_{-i}^{l-1}\}$, using the investment and contracting choices from the previous iteration, 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 (6), (9), and (10) for seller i and buyer j , using their new beliefs $\mathbf{Y}_{-i}^{l-1}$. (The superscript l refers to the fact that these are guesses for iteration l .)
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.3 Contracting and allocative efficiency: further analysis

Section 6.1 showed that LNG is allocated more efficiently in the short-run when no long-term contracts are used. This section investigates how allocative efficiency varies with the *extent* of long-term contracting.

The motivation behind this analysis is that the allocative efficiency effects of long-term contracts depend on how large sellers' contractual commitments are relative to their capacity. Recall that the key trade-off when it comes to allocative efficiency is that 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 there are potentially large gains from freeing up additional capacity for the spot market. As the contracted share of capacity decreases, though, the flexibility gain from freeing up capacity diminishes, so that at some point we would expect the market power effect to dominate.

Figure C10: Contracts and Allocative Efficiency



Note: The vertical axis plots aggregate industry surplus (in US\$ bn) from 2006 to 2017, discounted back to 2006. The horizontal axis shows the “cap” on the maximum share of capacity that can be contracted: for instance, a cap of 0.6 means that no more than 60% of each seller’s capacity can be contracted. To implement this, I proportionally scale down the contracts of any seller whose total contracted quantity exceeds the cap until the total quantity under contract equals the limit.

To investigate this, I consider counter-factual scenarios that place a cap on the maximum share of each seller’s capacity that can be contracted. For example, a cap of 0.6 means that no more than 60% of each seller’s capacity can be contracted. I investigate caps ranging from 0% (which corresponds to the scenario of no contracting) to 100% (which is the baseline scenario with unrestricted contracting). Consistent with the above discussion, I find a non-monotonic relationship between the level of contracting and allocative efficiency (Figure C10). Allocative efficiency is maximized when the cap on the share of capacity contracted is around 0.2: beyond that, further restrictions on contracting hurt allocative efficiency (due to the worsening distortions from market power).

C.4 Industry responses to demand shocks

In this section, I provide further details on the counter-factual simulation described in Section 6.2, where I explore the LNG industry response to a shutdown of Russian natural gas exports to Europe. I also discuss the results from a counter-factual exercise studying the impact of the Fukushima nuclear disaster in Japan.

Shutdown of Russian natural gas exports to Europe First, I consider the effect of a hypothetical shutdown of Russian natural gas exports to Europe. This issue rose to 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.⁷³

I consider two alternative contracting regimes for the LNG industry. The baseline regime involves rigid long-term contracts with resale restrictions, as we see in practice during the sample period. I compare this to a scenario where resale restrictions are prohibited in long-term contracts. For each of these counter-factuals, I numerically solve for the spot market equilibrium both with and without the demand shock.

As was discussed in Section 6.2, in the absence of resale restrictions, LNG sellers would respond by exporting considerably more LNG to Europe, leading to a substantially smaller increase in European spot prices in response to the shutdown of Russian natural gas exports.

Fukushima nuclear disaster Next, I explore the effect of the Fukushima nuclear disaster in Japan (in March 2011) on the LNG industry. In the months following the Fukushima disaster, Japan shut down all of its nuclear plants, which had previously accounted for around 25% of its total electricity generation capacity.⁷⁴ 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 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

⁷³The binding constraint is not so much spare regasification capacity, which is plentiful in European countries, but rather constraints in inter-connection capacity. This makes it challenging to deliver natural gas from European countries with LNG import capabilities (such as Spain, or France) to countries reliant who do not have their own LNG import facilities (such as Germany). See [IEA \(2022\)](#) for further discussion of this point.

⁷⁴Calculations based on 2010 data, which was obtained from BP's Statistical Review of World Energy, 2021.

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.

Once again I consider two alternative contracting regimes. The baseline takes the industry as it was, with rigid long-term contracts. I contrast this to a regime where no destination clauses or other resale restrictions are used in long-term contracts, as discussed in Section 6.2.

Table C14: Effect of Fukushima nuclear crisis, 2011 - 2013: with and without resale restrictions

	Benchmark	No Resale Restrictions
Δ Japanese LNG imports (mt)	29.3	46.6
Δ Imports from sellers with low contract share (mt)	6.3	6.4
Δ Imports from sellers with high contract share (mt)	23.0	40.2
Δ Japan Spot price (\$/MMBtu)	5.4	3.3
Δ Global Average Spot Price (\$/MMBtu)	2.4	2.6

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 table reports the impact of the Fukushima nuclear crisis on allocations and spot prices (averaged across the 2011 - 2013 period) under both these regimes. “ Δ European LNG imports” (first row) shows the increase in Japan’s LNG imports due to the Fukushima crisis under the two different contracting regimes; the second and third two rows decompose this increase in imports into imports coming from sellers with less than 50% of their capacity committed under long-term contracts (“low contract share”), and those with at least 50% their capacity committed under long-term contracts (“high contract share”). The final two rows show how the shutdown affects both the Japanese sales-weighted average spot price, as well as the sales-weighted global average spot price.

Table C14 illustrates the results. As we can see from comparing “Benchmark” with “No Resale Restrictions”, the industry responds more efficiently to the demand shock when long-term contracts do not feature any resale restrictions. In the absence of resale restrictions, sellers sell more LNG to Japan: in the baseline regime, Japan’s LNG imports rise by 29.3 million tonnes (mt) due to the Fukushima crisis, but would have risen even more (by 46.6 mt) if sellers were not bound by long-term contracts. This increase is primarily concentrated among sellers with more than 50% of their capacity committed to long-term contracts, consistent with the fact that highly contracted sellers are the least likely to be able to respond flexibly to demand shocks when they are inhibited by contractual rigidities. Spot prices are estimated to increase by \$5.4/MMBtu in Japan due to the Fukushima disaster, but would have only gone up by \$3.3/MMBtu absent resale restrictions.

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.

- BARTSCH, U. (1998): *Financial risks and rewards in LNG projects: Qatar, Oman, and Yemen*, Oxford Institute for Energy Studies.
- CAOUI, E. H. AND A. STECK (2023): “Network diversity, market entry, and the global internet backbone,” .
- GOMES, I. (2020): *The dilemma of gas importing and exporting countries*, 161, OIES Paper: NG.
- 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.
- 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.
- NATURAL GAS WORLD (2017): “Financing LNG in the New Era,” *Natural Gas World*, accessed: 2025-02-01.
- 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.
- 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 Shut-off to the European Union,” Working paper, International Monetary Fund.
- PIERRU, A., S. ROUSSANALY, AND J. SABATHIER (2013): “Capital structure in LNG infrastructures and gas pipelines projects: Empirical evidences and methodological issues,” *Energy policy*, 61, 285–291.
- SAMUELSON, P. A. (1952): “Spatial price equilibrium and linear programming,” *The American economic review*, 42, 283–303.
- SONGHURST, B. (2014): “LNG Plant Cost Escalation,” Oies paper, Oxford Institute for Energy Studies, accessed: 2024-10-17.
- 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 Appendix for “Long-term contracts and efficiency in the liquefied natural gas industry”

Nahim Bin Zahur

May 2025

S.1 Model Properties and Simulations

In this section, I explore further properties of the model developed in Section 4, with the help of two sets of Monte Carlo simulations. Firstly, I explore, both theoretically and through simulations, how bargaining power and outside options affects the contracting and investment decisions of the parties (Section S.1.1). Second, I explore how contracts affect short-run allocations, as well as the size and nature of contracting externalities (Section S.1.2).

S.1.1 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 then describe Monte Carlo simulations that illustrate them.

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:

$$K_i^{**} = \underset{K_i}{\operatorname{argmax}} \left[\underbrace{V_i^3(\mathbf{q}_i^{c,1}, \mathbf{q}_i^{c,3}, K_i, \mathbf{Y}_{-i}) + W_j^3(\mathbf{q}_i^{c,1}, \mathbf{q}_i^{c,3}, K_i, \mathbf{Y}_{-i}, \eta_{ij}^3)}_{\text{Sum of seller and buyer surplus}} - \Gamma_i(K_i, \mathbf{q}_i^{c,1}, \eta_i^2) \right] \quad (\text{S.11})$$

By contrast, the investment actually chosen by the seller maximizes the sum of the seller’s payoff from investing and the lump-sum transfer the seller receives from the buyer. Combining equations (7) and (8) yields the following equation for the seller’s choice of investment:

$$K_i^* = \underset{K_i}{\operatorname{argmax}} \left[\underbrace{\tau \left(V_i^3(\mathbf{q}_i^{c,1}, \mathbf{q}_i^{c,3}, K_i, \mathbf{Y}_{-i}) + W_j^3(\mathbf{q}_i^{c,1}, \mathbf{q}_i^{c,3}, K_i, \mathbf{Y}_{-i}, \eta_{ij}^3) \right)}_{\text{Sum of seller and buyer surplus}} - \Gamma_i(K_i, \mathbf{q}_i^{c,1}, \eta_i^2) + \right. \\ \left. \underbrace{(1 - \tau) V_i^3(\mathbf{q}_i^{c,1}, \mathbf{q}_{i,\setminus ij}^{c,3}, K_i, \mathbf{Y}_{-i})}_{\text{Seller's disagreement payoff}} - \underbrace{\tau W_j^3(\mathbf{q}_i^{c,1}, \mathbf{q}_{i,\setminus ij}^{c,3}, K_i, \mathbf{Y}_{-i})}_{\text{Buyer's disagreement payoff}} \right] \quad (\text{S.12})$$

Comparing equations (S.11) and (S.12) 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 . Then, 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 $[(V_i^3 + W_j^3) - \Gamma_i]$. 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. 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.

K_i^* will thus 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 Next, I report the results from Monte Carlo simulations to explore the relationship between bargaining power, investment and contracting. The simulations are based on a toy version of the model where a single seller interacts with a single buyer, and both the seller and the buyer have costly access to a spot market (that forms their “outside option” in case bargaining breaks down). The seller and buyer play the multi-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. In the toy model, I assume both the seller and the buyer must incur transaction costs (e.g., shipping costs) when trading on the spot market, which in turn determines their respective outside options.

Consistent with the theoretical predictions, I 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, as illustrated in Figure S.1. Similarly, the weaker the seller's outside option (as captured by a higher cost of selling to the spot market), the more severe the under-investment and the smaller the level of capacity the seller builds (Figure S.2).⁷⁵

Next, if I do allow the seller and buyer to sign ex-ante contracts, I find increased investment and welfare. Moreover, the lower the Nash bargaining weight of the seller, and 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 Figures S.3 and S.4). This is because the under-investment in Stage 2 is more severe when sellers have low bargaining leverage relative to the buyers; anticipating this, the seller and buyer sign larger ex-ante contracts.

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

Figure S.1: Effect of seller's Nash bargaining weight τ on K





Figure S.2: Effect of seller's cost of selling to spot market on K



⁷⁵I 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.

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

Figure S.3: Share of contracts signed in Stage 1 (ex-ante) vs. seller bargaining power





Figure S.4: Share of contracts signed in Stage 1 (ex-ante) vs. seller's cost of selling to spot market



S.1.2 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 S.1.2.1). In Section S.1.2.2, I then explore the conditions under which contracting externalities can emerge. For tractability, I do not endogenize the choice of contracts or investments, 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 ($d_{22} < d_{21}$). I assume that each seller i cannot produce beyond their capacity K_i and has a constant marginal cost of production for any output below K_i . 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 S.2):

1. Sim1, where firms have lots of capacity relative to demand and always produce below capacity.
2. Sim2, where firms have very limited capacity and will always produce at full capacity (even when they have no contracts).

Table S.1: 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 S.2: Simulations

	Sim1	Sim2
Description:	High capacity	Limited capacity
K_1	100	15
K_2	100	15

S.1.2.1 Allocative efficiency

To investigate how contracts affect allocations, I consider a variety of contract configurations (Table S.3). 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.

Simulation 1 (Sim1): firms always produce below capacity: Table S.4 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

Table S.3: Contract configurations

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

reducing the payoff they receive on their contracted output.

Table S.4: Effect of contracts when firms produce below capacity

		Prices			Trade flows					
	Contracts	Welfare	p_1	p_2	Seller 1		Seller 2		Total Prod.	Total 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 (ϵ_1, ϵ_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.

Simulation 2 (S2): firms always produce at capacity: Table S.5 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 S.5: Effect of contracts when firms produce at capacity

		Prices			Trade flows					
	Contracts	Welfare	p_1	p_2	Seller 1		Seller 2		Total Prod.	Total 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 (ϵ_1, ϵ_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 S.6 and Table S.7. 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 impede such reallocation. Table

S.6 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. (The results are similar when demand is low in market 1 but high in market 2). By contrast, as Table S.7 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 S.6: Effect of contracts when demand is high in market 1, low in market 2

	Contracts	Welfare	Prices		Trade flows				Total prod.	Total capacity
			p_1	p_2	Seller 1		Seller 2			
					q_{11}	q_{12}	q_{21}	q_{22}		
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 (ϵ_1, ϵ_2). I solve for spot prices and allocation, taking any contracted flows as given. The table presents averages across realizations of ϵ where $\epsilon_1 > 70$ and $\epsilon_2 < 50$.

Table S.7: Effect of contracts when demand is high in both market 1 and market 2

	Prices				Trade flows				Total prod.	Total capacity
	Contracts	Welfare			Seller 1		Seller 2			
			p_1	p_2	q_{11}	q_{12}	q_{21}	q_{22}		
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 (ϵ_1, ϵ_2). I solve for spot prices and allocation, taking any contracted flows as given. The table presents averages across realizations of ϵ where $\epsilon_1 > 70$ and $\epsilon_2 > 70$.

S.1.2.2 Contracting externalities

I next investigate whether contracts impose externalities on other firms. I take as the baseline case the Cournot 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 (Simulation 1), and when firms have limited capacity (Simulation 2).

Simulation 1 (Sim1): 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 S.8). 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 S.8: Effect of contract signed by seller 1, when firms produce below full capacity

	Contract quantity				Change in welfare, relative to baseline				
	q_{11}^c	q_{12}^c	q_{21}^c	q_{22}^c	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

Simulation 2 (Sim2): 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 S.9. 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 S.9 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 S.9 shows, when a buyer signs a contract, the average quantity that they purchase increases, and the average price they pay decreases.

Table S.9: Effect of contract signed by seller 1, when firms produce at capacity

	Contract quantity				Change in welfare, relative to baseline				
	q_{11}^c	q_{12}^c	q_{21}^c	q_{22}^c	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 S.9 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 S.9 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 S.9), 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 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 externalities imposed by contracts on sellers 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 S.9, 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.

S.2 Further Estimation Details and Descriptive Evidence

S.2.1 Contract Quantity Regressions

Table S.10 repeats the descriptive regression of contract quantity in Table 2, but with the inclusion of additional controls. Column (1) is identical to the first column of Table 2. In Column (2), I control 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. 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); the rule of law for importers is found to negatively predict the contract quantity. In Column (3), I consider a different way of measuring the bargaining leverage 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 bargaining leverage; with the opposite being the case if there are several other sellers signing contracts during the same year. However, the coefficients on these controls are small in magnitude and insignificant. Finally, Column (4) 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 decreases, consistent with the theoretical predictions.

S.2.2 Demand Estimates: Robustness Checks

I study how the choice of control variables and instruments affects the demand estimates (Table S.11). Column (1) repeats the baseline specification. In Column (2), I include, as an additional demand shifter, the importing country’s baseload electricity consumption (i.e., its electricity generation from baseload sources including nuclear, hydro and renewables). In Column (3), instead of including the country’s electricity consumption from fossil fuel and baseload sources separately, I include the country’s total electricity consumption as a control variable. Regardless of how exactly I include demand shifters, the resulting estimates of LNG demand (and average demand elasticities) are very similar.

I also investigate the effect of using different combinations of IVs. In addition to the three baseline instruments, I consider an additional IV for the spot price: total liquefaction capacity in the world in period t (abbreviated to “Global Liq. Cap.” in the tables). The greater the liquefaction 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 global LNG export capacity is uncorrelated with a country’s idiosyncratic demand shocks, after controlling for demand shifters. The logic behind the instrument is that LNG terminals take many years to build, and at the time the decision to invest is

Table S.10: Contract quantity regressions: include various controls

Dependent variable	Contract Quantity Share			
	(1)	(2)	(3)	(4)
Ex-ante*Relative distance, seller	0.17*** (0.044)	0.13*** (0.044)	0.16*** (0.045)	0.13*** (0.045)
Ex-ante*Relative distance, buyer	-0.15*** (0.037)	-0.12*** (0.037)	-0.15*** (0.038)	-0.13*** (0.039)
Relative distance, seller	-0.012 (0.017)	-0.023 (0.016)	-0.012 (0.017)	0.0078 (0.022)
Relative distance from buyer	0.011 (0.0091)	0.011 (0.0087)	0.010 (0.0090)	-0.0042 (0.012)
Ex-ante contract	0.024 (0.035)	0.029 (0.034)	-0.021 (0.059)	0.081** (0.036)
Contracted in the past		-0.042** (0.019)		
Rule of law, exporter		-0.0089 (0.0095)		
Rule of law, importer		-0.065*** (0.017)		
No. available buyers in same year			-0.0050 (0.0031)	
No. available sellers in same year			0.0011 (0.0050)	
Ex-ante*No. available buyers in same year			-0.0063 (0.0053)	
Ex-ante*No. available sellers in same year			0.016* (0.0085)	
Importer region fixed effects				Yes
Exporter region fixed effects				Yes
N	337	337	337	337
R ²	0.29	0.35	0.31	0.34

Note: Each observation is a long-term contract. The dependent variable is the contract quantity expressed as a share of plant capacity. The relative distance of an agent is their distance from alternative trading partners (i.e., median distance to all potential trading partners) divided by the distance from their current trading partner. Other controls include the distance between the seller and buyer, the logarithm of plant capacity, an indicator for greenfield plants, an indicator for extensions, a time trend, and (in Column (4)) importer & exporter region fixed effects.

made, it is difficult to foresee idiosyncratic demand shocks that are realized several years later.

Columns (4) and (5) in Table S.11 presents demand estimates from different combinations of these four instruments. After including the liquefaction capacity instrument, the demand curve is estimated to be somewhat more elastic than in the baseline specification, though the elasticities are still broadly in the same range (-1.27 as opposed to -0.83 in the baseline).

Table S.11: Demand Estimates: Robustness Checks

	(1)	(2)	(3)	(4)	(5)
Spot Price (USD/MMBtu)	-0.106 (0.055)	-0.090 (0.052)	-0.106 (0.051)	-0.160 (0.042)	-0.161 (0.042)
Log (Elec. Cons., fossil)	1.115 (0.201)	1.274 (0.175)		1.130 (0.212)	1.130 (0.212)
Min. Temp. (standardized)	-0.092 (0.027)	-0.107 (0.025)	-0.096 (0.029)	-0.104 (0.029)	-0.104 (0.029)
Oil Price (USD/barrel)	0.011 (0.006)	0.011 (0.005)	0.012 (0.005)	0.017 (0.005)	0.017 (0.005)
Log (Elec. Gen., baseload)		0.800 (0.107)			
Log (Elec. Cons.)			2.331 (0.285)		
Kleibergen-Paap F-stat	9.83	9.02	10.36	19.21	14.75
Olea-Pflueger F-stat	10.08	9.10	10.47	17.31	13.11
Median elasticity	-0.83	-0.71	-0.84	-1.26	-1.27

Note: Each observation is an importer-year-quarter pair. There are 815 observations in all specifications except for (2) (which has 790 observations). The dependent variable is the logarithm of total LNG imports in country j in year-quarter t . In Column (1), the control variables, fixed effects and instruments are the same as in Table B8. Column (2) also controls for the country's total electricity generation from baseload sources ("Log (Elec. Gen., baseload)"). In Column (3), I control for the country's total electricity consumption from both baseload and fossil fuel sources ("Log (Elec. Cons.)"). In Column (4), I replace the rival electricity generation from baseload IV with a different IV, the total liquefaction capacity in the world in period t ("Global Liq. Cap."). Column (5) includes all four instruments. All regressions include importing country and quarter fixed effects. Robust standard errors are in parentheses.

S.2.3 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 32). 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 $u_1, \dots, u_{L_s}$, and each buyer's expected payoffs can be approximated by a set of L_b basis functions $\phi_1, \dots, \phi_{L_b}$:

$$\pi_{it}^s(\mathbf{q}_t^c, \mathbf{K}_t) \simeq \sum_{l=1}^{L_s} b_l^s u_l(\mathbf{q}_t^c, \mathbf{K}_t, \mathbf{x}_t) \quad (\text{S.23})$$

$$\pi_{jt}^b(\mathbf{q}_t^c, \mathbf{K}_t) \simeq \sum_{l=1}^{L_b} b_l^b \phi_l(\mathbf{q}_t^c, \mathbf{K}_t, \mathbf{x}_t) \quad (\text{S.24})$$

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

In order to estimate the approximating parameters, I first randomly draw S combinations of $\mathbf{q}_t^c$ (vector of contracted quantities) and $\mathbf{K}_t$ (capacities). For each of these S draws of the states, I randomly draw D realizations of $\boldsymbol{\varepsilon}_t$ (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 $\mathbf{q}_t^c, \mathbf{K}_t, \mathbf{x}_t$, as well as $\pi_{it}^s(\mathbf{q}_t^c, \mathbf{K}_t, \mathbf{x}_t)$ for each seller i and $\pi_{jt}^b(\mathbf{q}_t^c, \mathbf{K}_t, \mathbf{x}_t)$ for each buyer j . I regress $\pi_{it}^s(\mathbf{q}_t^c, \mathbf{K}_t, \mathbf{x}_t)$ on basis functions of $(\mathbf{q}_t^c, \mathbf{K}_t, \mathbf{x}_t)$ in order to obtain the seller approximating parameters b_i^s ; and regress $\pi_{jt}^b(\mathbf{q}_t^c, \mathbf{K}_t, \mathbf{x}_t)$ on basis functions of $(\mathbf{q}_t^c, \mathbf{K}_t, \mathbf{x}_t)$ in order to obtain the buyer approximating parameters, b_j^b .

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).⁷⁶ 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.

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, and instead, only keep track of two sufficient statistics: 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 equilibria in dynamic models.⁷⁷

With the above assumption, I am able to reduce the dimension of the *endogenous* state variables (i.e., capacity $\mathbf{K}_t$ and contract quantities $\mathbf{q}_t$). However, the dimension of $\mathbf{x}_t$ (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 $\mathbf{x}_t$ 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 $\mathbf{x}_t$, and instead only keep track of selected indices that succinctly summarize the effect of $\mathbf{x}_t$ on their payoffs. Define $z_{jt} = \alpha + x_{jt} \theta_{dx} + \theta_j + \theta_{quarter}$. z_{jt} can be thought of buyer j 's demand state: as z_{jt} increases, the buyer's demand curve for LNG shifts outwards, which will change buyers' consumer surplus. 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 $z_j = \sum_t z_{jt}$,

⁷⁶I 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.

⁷⁷Empirical applications of this approach include Sweeting (2015), Chen and Xu (2021), Gerarden (2023), and Jeon (2022).

which is a measure of aggregate demand for LNG. 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.

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 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 to depend 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, when approximating seller’s payoff, I include $\bar{d}_{it}^z$, 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_{jt}^c , the sum of rival contracted quantities (Q_{-jt}^c), total capacity in that year K_t , 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_{jt}^c , Q_{-jt}^c , and K_t , with interactions of those terms with z_{jt} and $\bar{d}_j$.⁷⁸ The approximation parameters are estimated by an OLS regression of buyer expected payoffs on these selected basis functions, using the simulated sample.⁷⁹

Seller payoff approximation: Sellers’ expected payoffs are functions of their total contracted quantity Q_{it}^c , their capacity K_{it} , the sum of rival contracted quantities (Q_{-it}^c) and rival capacities in that year (K_{-it}), as well as aggregate demand z_t and the seller’s demand-weighted average distance from other buyers $\bar{d}_{it}^z$. As with buyers, I assume sellers’ payoffs include a quadratic function of the endogenous variables as well as interaction terms between each of the endogenous variables and the two indices summarizing the exogenous variables.

But the quadratic basis functions, by themselves, do not capture capacity constraints all that well, which are crucial for approximating seller payoffs. To better capture capacity constraints, I also include several higher-order polynomial terms in Q_{it}^c and K_{it} : $Q_{it}^{c^2} * K_{it}$, $Q_{it}^c * K_{it}^2$, K_{it}^3 and $Q_{it}^{c^3}$. Finally, I also include a basis function especially designed to capture capacity constraints, $Q_{it}^{c^3} / K_{it}^2$. The inclusion of this basis function implies that the derivative of the seller’s expected payoff with respect to Q_{it}^c is a quadratic function of Q_{it}^c / K_{it} (or the share of the seller’s capacity that is committed under long-term contracts). The idea behind this is that as Q_{it}^c gets larger (relative to K_{it}), the seller’s capacity utilization is higher and the seller is more likely to be highly capacity-

⁷⁸I exclude interactions between z_{jt} and contract quantities (Q_{jt}^c , Q_{-jt}^c), since I found that including those interaction terms resulted in poorly behaved approximations.

⁷⁹These results are available upon request.

constrained, which will lower their expected payoffs (since they will be unable to take advantage of periods when spot prices are high). I found that these particular basis functions yielded sensible and well-behaved approximations of seller payoffs; experimentation with many other combinations of basis functions yielded largely similar results.⁸⁰

S.2.4 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 (which is available for some countries) provides information on the annual LNG prices paid by an importer for LNG imported from different exporting countries. Using these 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 Eikon, 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 linear equations of the following form:⁸¹

$$p_{ijt}^c = a_{ij} + b_{ij} \sum_{s=0}^T w_s p_{t-s}^{oil} \quad (\text{S.25})$$

where p_{ijt}^c is the per-unit price of LNG that buyer j pays to seller i in period t (which is available

⁸⁰The approximation parameters are estimated by an OLS regression of seller expected payoffs on these selected basis functions, using the simulated sample. The coefficients and regression results are available upon request.

⁸¹Oil 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)

from the customs data). $\sum_{s=0}^T w_s p_{t-s}^{oil}$ 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. Though a_{ij} and b_{ij} are unknown, they can be estimated by regressing the LNG price on lags of the oil price, once we have enough observations of p_{ijt}^c for the same exporter-importer pair i, j .

This approach of inferring contract pricing formulae from customs data has been used by [Agerton \(2017\)](#) and is also commonly used by industry analysts (see, for example, [Flower and Liao, 2012](#)). There are, however, some limitations. First, customs data does not differentiate between long-term contracts and spot; as such, if importer j buys from exporter i using both long-term contracts and spot trade, the price p_{ijt}^c reported in the customs data will be a share-weighted average of the contract and spot price. Second, an importer j and an exporter i may have multiple long-term agreements in force at any given time; in that case, the price p_{ijt}^c 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 use “S-curves” where they follow the linear pricing formula (S.25) 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.⁸² 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 when structurally estimating the model).

Equation (S.25) can be rewritten as follows, where $b_{ij}^s = b_{ij} w_s$:

$$p_{ijt}^c = a_{ij} + \sum_{s=0}^{t-T} b_{ij}^s p_{t-s}^{oil} \quad (\text{S.26})$$

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.⁸³

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 contract prices in-

⁸²See, for example, [Flower and Liao \(2012\)](#), for more discussion on the use of S-curves in Asian LNG contracts.


⁸³The regression results are available upon request.

ferred from customs data. The raw prices reported in the customs data are truncated for most contracts.⁸⁴ As such, I use the estimated contract price formulas (based on my empirical estimates of equation (S.26)) together with information on oil prices to construct the long-term average contract price across the lifetime of each contract.⁸⁵ I then aggregate these contract-level average prices back to the exporter-importer level, using the contract quantities as the weights, to calculate the average contract price p_{ij}^c 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_{ij}^c (the lump-sum transfer) for every contract signed by importer j and exporter i , as well as the contract quantity they agree to (q_{ij}^c) and the contract duration. Together, these are used to calculate the per-unit price for every contract. I 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, $\hat{p}_{ij}^c$.

Figure S.5 plots the model-predicted average contract price on the x-axis ($\hat{p}_{ij}^c$) and the average contract price inferred from customs data on the y-axis (p_{ij}^c). The two are fairly highly correlated, with a correlation coefficient of 0.478, suggesting the model provides reasonable estimates of contract prices.

Figure S.5: 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 S.2.4 for details on how contract prices are computed from the customs data.

⁸⁴Most 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.

⁸⁵This requires forming an expectation of future oil prices. I assume firms use a simple AR(1) model to forecast future prices.

S.2.5 Back-of-the-envelope calculation of the impact of ex-ante contracts on the average cost of investment

A benefit of ex-ante contracts is they may enable sellers to obtain more debt (by lowering the risk profile of the project) and therefore lower the cost of financing the project, given that debt is generally cheaper than alternative forms of financing (such as equity). In this section, I discuss a simple back-of-the-envelope calculation of this benefit, using existing industry sources. This is useful as a way to benchmark the estimate of the financing benefit of ex-ante contracts that I obtain from the structural model.

I carry out this calculation in two steps. First, I draw on [Bartsch \(1998\)](#), who provides an estimate of the average reduction in investment costs from using debt to finance LNG projects. Second, I determine the contribution of ex-ante contracts to the debt share of LNG projects, relying on [Ruester \(2015\)](#). Combining these two steps, I arrive at a rough back-of-the-envelope estimate of the impact of ex-ante contracts on investment costs.

[Bartsch \(1998\)](#) carries out a detailed financial analysis of several LNG projects in the Middle East (in Qatar, Yemen and Oman), and finds that the use of debt lowers the average investment cost by 10-12%. This provides an upper bound on the financing benefit from ex-ante contracts: if it was impossible to secure any debt without ex-ante contracts, then the financing benefit from using ex-ante contracts would amount to 10-12%.

In practice, LNG projects may be able to obtain some debt even if they do not use ex-ante contracts. The relationship between ex-ante contracting and the seller's ability to secure debt is empirically studied by [Ruester \(2015\)](#), who finds (using a regression analysis) that increasing the share of capacity committed under ex-ante long-term contracts from 0% to 100% increases the debt share (i.e., the share of the funding coming from debt) of a project by 0.257.

In my dataset, the average share of capacity committed under ex-ante long-term contracts is 61.3%, implying (based on [Ruester \(2015\)](#)'s estimates) that ex-ante contracts increase the debt share by $0.257 \times 61.3\% = 15.7\%$. Given that 40% of all funding for liquefaction projects comes from project finance, and the debt share for such projects is around two-thirds on average, the average debt share for LNG projects is around 27%. This implies that the use of ex-ante contracts accounts for around 60% of the overall debt taken on by the project. Since debt lowers investment costs by 10-12%, the use of ex-ante contracts translates to a 6-7% reduction in investment costs – quite close to the structural model estimate of 5.1%.

Additional References

Chen, Yanyou and Daniel Yi Xu, “A Structural Empirical Model of R&D, Firm Heterogeneity, and Industry Evolution,” *Journal of Industrial Economics*, 2021.

- Flower, Andy and Jane Liao**, “LNG Pricing in Asia,” in Jonathan P Stern, Hongyan Wang, and Tianxiao Fan, eds., *The pricing of internationally traded gas*, Oxford: Oxford University Press Oxford Institute For Energy Studies, Oxford, 2012, chapter 11.
- Gerarden, Todd D**, “Demanding innovation: The impact of consumer subsidies on solar panel production costs,” *Management Science*, 2023, 69 (12), 7799–7820.
- Hendel, Igal and Aviv Nevo**, “Measuring the Implications of Sales and Consumer Inventory Behavior,” *Econometrica*, 2006, 74, 1637–73.
- Jeon, Jihye**, “Learning and investment under demand uncertainty in container shipping,” *The RAND Journal of Economics*, 2022, 53 (1), 226–259.
- Nevo, Aviv and Federico Rossi**, “An Approach for Extending Dynamic Models to Settings with Multi-Product Firms,” *Economics Letters*, 2008, 100, 49–52.
- Sweeting, Andrew**, “A model of non-stationary dynamic price competition with an application to platform design,” *Available at SSRN 2672027*, 2015.