

ESSAYS ON INTERNATIONAL ENERGY AND TRANSPORTATION MARKETS

A Dissertation

Presented to the Faculty of the Graduate School

of Cornell University

in Partial Fulfillment of the Requirements for the Degree of

Doctor of Philosophy

by

Nahim Bin Zahur

May 2020

© 2020 Nahim Bin Zahur
ALL RIGHTS RESERVED

ESSAYS ON INTERNATIONAL ENERGY AND TRANSPORTATION MARKETS

Nahim Bin Zahur, Ph.D.

Cornell University 2020

This dissertation consists of three separate essays that seek to examine the interaction between firm behavior and government policy in international energy and transportation markets. Chapters 1 and 2 are closely related in that they both focus on the liquefied natural gas (LNG) industry, one of the fastest-growing energy markets globally. The first chapter measures market power in the LNG spot market and studies how market power influences pricing, trade and welfare. I develop a method for inferring market conduct that utilizes information on sellers' pricing and quantity decisions across multiple geographically segmented markets. My test for market conduct is based on the observation that sellers exercising market power engage in third-degree price discrimination, whereas sellers behaving competitively do not. Using data from 2006 to 2017 on spot market trade flows, spot prices, shipping costs and seller capacities, I estimate a structural model of LNG trade and pricing that incorporates spatial differentiation, capacity constraints and trade frictions and flexibly nests different models of seller market power. I find that seller decisions are consistent with a Cournot model and unlikely to be generated by a competitive model. The total deadweight loss from market power is estimated to be USD 12 billion, or about 4.5% of total revenue. I find that market power plays a key role in exacerbating inter-regional price differentials.

Perhaps the most distinct feature of the LNG industry is its reliance on long-term contracts (typically 20 years in length). Long-term contracts facilitate efficient levels of investment by sellers faced with the risk of ex-post holdup. Contractual rigidities, however, reduce the ability of the market to respond flexibly to demand uncertainty.

In the second chapter, I provide the first empirical analysis of the trade-off between hold-up risk and contract rigidity, using the liquefied natural gas (LNG) industry as a case study. I structurally estimate a model of contracting, investment and spot trade that incorporates hold-up risk and contractual rigidities, using a rich dataset on LNG contracts, investment, trade flows, and spot prices. LNG buyers are estimated to have substantial bargaining power, implying that sellers face considerable hold-up risk when making investment decisions. In the short-run, allocative efficiency would improve through reduced use of long-term contracts. However, investment decreases by 35% in the absence of long-run contracting, suggesting that long-term contracts play a significant role in combatting holdup at the cost of short-term allocative efficiency.

The final chapter of this dissertation is co-authored with Panle Jia Barwick and Myrto Kalouptsi. Despite the historic prevalence of industrial policy and its current popularity, few empirical studies directly evaluate its welfare consequences. We examine an important industrial policy in China in the 2000s, aiming to propel the country's shipbuilding industry to the largest globally. Using comprehensive data on shipyards worldwide and a dynamic model of firm entry, exit, investment, and production, we find that the scale of the policy was massive and boosted China's domestic investment, entry, and world market share dramatically. On the other hand, it created sizable distortions and led to increased industry fragmentation and idleness. The effectiveness of different policy instruments is mixed: production and investment subsidies can be justified by market share considerations, but entry subsidies are wasteful and dissipated by the entry of inefficient firms. Finally, how the policy is implemented matters: the distortions could have been significantly reduced by implementing counter-cyclical policies and by targeting productive firms.

BIOGRAPHICAL SKETCH

Nahim Bin Zahur joined the doctoral program in the Department of Economics at Cornell University in the fall of 2014. His research interests are in empirical industrial organization and energy and environmental economics. A native of Bangladesh, he received a BA in Economics from the National University of Singapore in 2011. Prior to joining Cornell, he spent 3 years as an Analyst (and later Research Associate) at the Energy Studies Institute in Singapore.

Dedicated to my wonderful parents, Zahurul Islam and Fahmida Akhand.

ACKNOWLEDGEMENTS

My deepest gratitude goes to my committee members and mentors Panle Jia Barwick, Jean-Francois Houde, Shanjun Li, and Giulia Brancaccio for their invaluable advice and guidance. I have truly learnt a lot from them over the last few years, and they have each played a crucial role in my development as an economist. I am especially grateful to Panle Jia Barwick, my committee chair, for guiding me through all stages of my dissertation, and for her constant encouragement. I am indebted to Jean-Francois Houde, without whose insightful advice and support I could not have written this dissertation.

My first brush with economics research came while working with Sajjad Zohir. His encouragement and mentorship was pivotal in my decision to become an economist, and for this I thank him. I am thankful to Parkash Chander, who guided me in writing my undergraduate thesis, and to Tilak Doshi, who sparked my enduring interest in energy economics.

In the course of writing this dissertation, I have greatly benefitted from the generous advice and help provided by numerous individuals, both at Cornell and elsewhere. I am especially thankful to Myrto Kalouptsi for her dual role as a mentor and co-author. I sincerely thank Todd Gerarden, Ivan Rudik, Francesca Molinari, Jorg Stoye, Penny Sanders, and Seth Sanders for several useful discussions. I learnt a great deal from discussions with my peers and friends at Cornell, including Abhishek Ananth, Lucy Xiaolu Wang, Deyu Rao, Esteban Mendez-Chacon Cambiasso, Yu She, Tilahun Emiru, Armand Sim, Jee-Hun Choi, and Qi Wu.

I thank Victor Aguirregabiria, Matt Backus, Levon Barseghyan, Vivek Bhattacharya, Diego Cardoso, Maura Coughlin, Thomas Eisenberg, Paul Grieco, Jerry Hausman, Christopher Huckfeldt, Ian Lange, Richard Lowery, Luca Maini, John McLaren, Joris Pinske, James Roberts, Mark Roberts, Karl Schurter, Jesse Shapiro, Paolo Somaini, Chris Sullivan, Mathieu Taschereau-Dumouchel, Matthew Thirkettle, Daniel Xu, Yang

Zhang, and seminar participants at the Northeast Workshop on Energy Policy and Environmental Economics, Western Economic Association International Conference, Energy Research Insights for Decisionmaking Conference, Cornell IO Workshop, Cornell Dyson Energy and Environmental Economics Workshop, and Penn State University for helpful comments. Generous funding for my dissertation was provided by the Alfred P. Sloan Foundation Pre-doctoral Fellowship on Energy Economics, awarded through the National Bureau of Economic Research (NBER), as well as the C. V. Starr Fellowship in East Asian Studies offered by Cornell's East Asia Program.

No words can suffice to thank my parents, Zahurul Islam and Fahmida Akhand, my brother Fahim Bin Zahur, and my sister Nabila Binte Zahur. Their love and support has been my constant source of strength.

TABLE OF CONTENTS

Biographical Sketch	iii
Dedication	iv
Acknowledgements	v
Table of Contents	vii
List of Tables	ix
List of Figures	xi
 1 Market Power and Spatial Price Discrimination in the Liquefied Natural Gas Industry	 1
1.1 Introduction	1
1.2 Industry and Data	9
1.3 Descriptive Evidence of Market Power	18
1.4 Model	28
1.4.1 Model Description	29
1.4.2 Equilibrium	34
1.5 Estimation	38
1.5.1 Estimation of Demand Curve	38
1.5.2 Estimation of Cost and Conduct Parameters	41
1.6 Conclusion	48
 2 Long-term Contracts and Efficiency in the Liquefied Natural Gas Industry	 50
2.1 Introduction	50
2.2 Industry and Data	58
2.3 Descriptive Evidence	68
2.4 Model	76
2.4.1 Overview of model	76
2.4.2 Demand, production and spot trade	78
2.4.3 Contracting and Investment	84
2.4.4 Equilibrium properties of the contracting and investment game	95
2.5 Estimation	101
2.5.1 Estimation of Demand Curve	101
2.5.2 Estimation of Cost Parameters	105
2.5.3 Estimation of contracting and investment model	107
2.6 Counter-factual Analysis	116
2.7 Conclusion	118
 3 China's Industrial Policy: an Empirical Evaluation	 119
3.1 Introduction	119
3.2 Industry Background and Data	126
3.2.1 Industry Background	126
3.2.2 Data	129
3.2.3 Descriptive Evidence and Summary Statistics	131

3.3	Model	134
3.4	Estimation Strategy	143
3.4.1	Estimation of Static Parameters	144
3.4.2	Estimation of Dynamic Parameters	147
3.5	Results	154
3.5.1	Static Demand and Production Cost Estimates	154
3.5.2	Dynamic Parameters: First Stage	159
3.5.3	Dynamic Parameters: Second Stage	160
3.6	Evaluation of China's Industrial Policies	163
3.6.1	Evaluation of Industrial Policy	164
3.6.2	Evaluation of Consolidation Policies	174
3.6.3	Rationales for Industrial Policy	176
3.7	Conclusion	179
A	Appendices for Chapter 1	180
A.1	Monte Carlo Simulations	180
A.1.1	Assumptions	180
A.2	Derivation of first-order condition in the case of linear demand and quadratic costs	185
B	Appendices for Chapter 2	187
B.1	Long-term contracts and short-run allocations: Monte Carlo simulations	187
C	Appendices for Chapter 3	192
C.1	Estimation Details	192
C.1.1	Additional Charts and Tables	192
C.1.2	Calibrating the Fixed Cost	198
C.1.3	Production cost estimates: robustness exercises	199
C.1.4	Estimating the Investment Policy Function	204
C.1.5	First-stage Policy Functions and State Transition Estimates . . .	206
C.1.6	Value function approximation	209
C.2	Implementation of Counterfactual Analyses	211
C.3	A Simple Model of Subsidies	212

LIST OF TABLES

1.1	Summary Statistics	18
1.2	Reduced Form Evidence for Market Power: Regression of Share of Spot Sales to market j on Price at market j	26
1.3	Demand Curve Estimates	39
1.4	Conduct and Cost Parameter Estimates	45
2.1	Summary Statistics	68
2.2	Regression of share of contract volume signed before final investment decision date on characteristics of the export project	71
2.3	Reduced Form Evidence for Contracting Frictions: Regression of Share of LNG Sales to market j on Price at market j	75
2.4	Demand Curve Estimates	102
2.5	Demand estimation: adding other demand shifters	104
2.6	Conduct and Cost Parameter Estimates	106
2.7	Parameter estimates using minimum distance estimator	113
2.8	Average capacity with and without contracting	117
2.9	Decomposing capacity reduction by type of seller	117
3.1	Summary Statistics	133
3.2	Cost function estimates	157
3.3	Estimates of investment cost and scrap value parameters	161
3.4	Entry Cost Distribution (Mean), billion RMB	162
3.5	Comparison of Different Subsidies	169
3.6	Effect of subsidies on productive and unproductive firms	171
3.7	Pro-cyclical vs. counter-cyclical industrial policy	173
A.1	Monte Carlo Assumptions	181
A.2	Comparison of competitive and Cournot outcome,	183
A.3	Scenario with moderate demand in both markets	183
A.4	Scenario with high demand in market 1, low demand in market 2 . . .	183
B.1	Assumptions maintained in all scenarios	188
B.2	Simulations	188
B.3	Contract configurations	188
B.4	Effect of contracts when firms produce below capacity	189
B.5	Effect of contracts when firms produce at capacity	190
B.6	Effect of contracts when demand is high in market 1, low in market 2 .	191
B.7	Effect of contracts when demand is high in both market 1 and market 2	191
C.1	Shipbuilding National Industrial Policies	192
C.2	Demand estimates	193
C.3	Actual vs. Simulated Exit	194
C.4	Actual vs. Simulated Entrants	194

C.5	Impact of Subsidies on Ship Prices	194
C.6	Comparing Production and Investment Subsidies	195
C.7	Targeting subsidies to White List firms	195
C.8	Cost function estimates: perfect competition vs. Cournot competition .	200
C.9	Cost function estimates, pooling data across China/Japan/South Korea .	201
C.10	Cost function estimates: learning	202
C.11	Cost function estimates: additional robustness checks	203
C.12	AR(1) estimates for state transition processes	206
C.13	Estimates of exit policy function	207
C.14	Estimates of the investment policy function	208

LIST OF FIGURES

1.1	Histogram of long-term contract durations	13
1.2	Share of short-term contracts and spot transactions in LNG trade	13
1.3	Liquefaction capacity, regasification capacity and LNG trade over time	14
1.4	Industry concentration (four-firm concentration ratio) over time	19
1.5	Shipping costs (\$/MMBtu) for major LNG routes	19
1.6	LNG Spot Prices in Different Regions	23
1.7	Short-run and spot LNG imports to different regions	23
1.8	Relationship between total spot sales and number of countries that a seller exports to	28
1.9	Actual vs. Predicted Spot Prices	46
1.10	Price Differential between Northeast Asia and Southwestern Europe: Baseline vs. Competitive	48
2.1	Histogram of long-term contract durations	63
2.2	Share of short-term contracts and spot transactions in LNG trade	63
2.3	Liquefaction capacity, regasification capacity and LNG trade over time	64
2.4	Gap between contracting year-month and investment year-month	69
2.5	Cumulative share of capacity signed under long-term contracts	69
2.6	LNG Spot Prices in Different Regions	72
2.7	LNG imports under long-term contracts to different regions	72
2.8	Short-run and spot LNG imports to different regions	72
2.9	Stages of the game	86
2.10	1 seller, 1 buyer, no Stage 1 contracting: effect of seller bargaining power on K	97
2.11	Share of contracts signed in Stage 1 (pre-investment) vs. seller bargaining power	98
2.12	Actual vs. Predicted Spot Prices in Japan	107
2.13	Estimated cost of investment vs. reported cost of investment	115
3.1	China's Market Share Expansion	132
3.2	Entry of New Shipyards	132
3.3	Investment by Chinese Shipyards	132
3.4	Ship Prices	132
3.5	No of firms, with and without subsidies	167
3.6	Investment, with and without subsidies	167
3.7	Average firm cost-efficiency $\bar{s}_{jt}$ with subsidies during the boom vs. subsidies during the bust	174
C.1	China's Expansion in Major Industries	196
C.2	Simulated vs. actual investment	196
C.3	HHI for Chinese Shipbuilding, with and without subsidies	197
C.4	Industry Profits Under Different White Lists	197

CHAPTER 1

**MARKET POWER AND SPATIAL PRICE DISCRIMINATION IN THE
LIQUEFIED NATURAL GAS INDUSTRY**

Nahim Bin Zahur

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

1.1 Introduction

The modern history of the liquefied natural gas (LNG) industry has been characterized by periods when spot prices differ sharply across markets. A case in point is the

Fukushima nuclear disaster in March 2011, which prompted a shutdown of Japan's nuclear reactors and a sharp increase in Japan's LNG demand. During the year preceding the Fukushima disaster, Asian LNG spot prices were \$1/MMBtu higher than European spot prices. During the three years after the disaster, though, Asian prices were on average \$4.75/MMBtu (or 40%) higher than European prices.

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

The higher prices paid by buyers in Asia compared to buyers in Europe for the same commodity has sparked a debate about whether LNG suppliers are engaging in price discrimination. It has been common for LNG buyers in Asia to refer to an "Asian price premium" (see, for instance, Koyama, 2012). In response to high Asian prices and the perceived premium paid by Asian buyers, there have been several attempts by Asian importing countries to form buyer's alliances.¹ Yet the existence of pricing differentials across regions is insufficient by itself to establish that sellers exercise market power. The industry generally operates at very high levels of capacity utilization, demand

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

is fairly inelastic, and shipping costs are important. Under these conditions, prices may differ across markets even under competitive behavior. Finally, though the LNG industry is concentrated and arbitrage is costly, LNG suppliers often have strong ties with governments and it is by no means a given that they would exercise market power in a way consistent with standard oligopoly models. Whether or not there exists market power in LNG is therefore an empirical question.

In this paper, I ask whether there is evidence that sellers exercise market power in the LNG spot market. The goals of the paper are to measure market power in the LNG spot market, and study how market power influences pricing, trade and allocative efficiency. My method for inferring market conduct utilizes information on both the pricing and quantity decisions of sellers across multiple geographically segmented markets. I build on the observation that sellers exercising market power engage in third-degree price discrimination by charging different prices (net of transportation costs) to different buyers, whereas sellers behaving competitively charge the same prices to all buyers. Moreover, in the presence of capacity constraints, discriminatory behavior on the part of sellers will distort the allocation of goods across buyers.

The key challenge in identifying market power is that in the presence of capacity constraints and shipping costs, a competitive model can also generate price dispersion across markets. Consider an industry with geographically differentiated buyers and sellers. Even in the absence of capacity constraints, competitive prices paid by different buyers will differ due to shipping costs: a buyer located far away from sellers will generally pay a higher price than a buyer with a more favorable location. If in addition sellers are capacity constrained, the net price received by sellers (i.e. price minus shipping costs) will also not be equalized across destinations, because a capacity-constrained seller will not sell to all possible destinations. Thus without any quantity

information, price differentials, whether measured at the buyer or seller level, cannot be directly interpreted as evidence of market power. It is the combination of price differentials and quantity allocation decisions that indicate potential violations of the competitive model. If the net price a seller receives differs within the set of buyers to which they sell positive quantities, the seller is unlikely to be acting competitively. Finally, the existence of unobserved trade costs and other frictions (such as unobserved trade costs) can also result in deviations from competitive behavior, and the analysis needs to distinguish between market power and other trade frictions.

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

Although the pattern of sellers selling to multiple markets at different net prices suggests deviations from perfectly competitive behavior, it does not answer the question of whether these deviations are the result of market power or some other trade frictions. I thus carry out a reduced-form test for whether sellers exercise market power that is designed to tease apart the role of market power from other frictions. The intuition behind the test is that under the hypothesis of market power, large sellers have a weaker incentive to increase sales to a market with high demand than small sellers, since by

withholding sales from that market they can push up the price and increase their total revenue from that market. As such, large sellers will respond less sharply to price changes than small sellers. Under perfect competition subject to trade frictions, by contrast, sellers ought to be equally responsive to price changes irrespective of their total sales. I find that spot sales of small sellers are indeed more responsive to price differentials than the spot sales of large sellers, consistent with the existence of market power in the LNG spot market.

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

I estimate this model using data on capacity, trade flows, spot prices and shipping

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

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

My paper contributes to a growing literature detecting market power and studying the welfare consequences of market power in energy industries. A rich literature analyzes market power in restructured electricity markets in California (Borenstein *et al.*, 2002; Puller, 2007) and elsewhere (Wolak, 2000; Fabra and Toro, 2005; Sweeting, 2007; Hortacsu and Puller, 2008). The role of OPEC in the oil market has long been the

subject of academic study, with a recent example being Asker *et al.* (2019). In contrast, there has been much less work on market power in the LNG industry. The only existing paper that I am aware of is Ritz (2014), who develops a theoretical model to illustrate how LNG exporters can price discriminate across markets and argues that market power can rationalize observed inter-regional price differentials. My paper, in contrast, utilizes panel data on LNG trade flows and spot prices to structurally estimate a model of the LNG market. This allows me to both measure the level of market power and analyze the extent to which observed price differentials can be explained by market power, while accounting for the role played by capacity constraints, shipping costs and trade frictions.

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

Finally, my paper is related to the empirical literature documenting price discrimination and studying the effects of price discrimination in a variety of settings

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

The analysis in this paper focuses exclusively on the spot market in LNG, taking as given both long-term contracts and export capacity. In the next chapter of my dissertation, I endogenize both long-term contracts and investment and study the link between long-term contracts and efficiency. A key rationale behind long-term contracts is that they facilitate efficient levels of irreversible investments made by sellers faced with hold-up risk (Williamson, 1979). This is likely to be very relevant in the LNG industry, where an exporter must construct an export terminal that typically costs several billion dollars and at least 4 years to build before they can engage in LNG exports. However contracts can result in short-run inefficiency since they do not respond flexibly in response to demand shocks. In the next chapter, I study the trade-off between hold-up risk (which is improved by contracts) and short-run misallocation (which is worsened by the use of contracts).

The remainder of this paper is divided as follows. Section 3.2 discusses key institutional features of the LNG industry and describes the dataset I use for the analysis. Section 2.3 provides descriptive evidence for market power, including a reduced-form test for whether sellers exercise market power. Section 2.4 develops a model of trade and pricing in this industry. Section 2.5 describes estimation of the key parameters of the model that characterize demand, cost and conduct. Section 2.7 concludes.

1.2 Industry and Data

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

Many countries with limited gas reserves rely on imported liquefied natural gas (LNG) to meet their natural gas needs. For countries such as Japan and Korea that do not have their own gas reserves and cannot import gas via pipelines, LNG provides their

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

only source of natural gas. Other countries, such as China and Spain, import natural gas via both pipelines and LNG. The largest producer of LNG is Qatar, followed by Australia, Malaysia and Indonesia. Japan is the largest importer of LNG, followed by China, Korea, India, Taiwan and Spain.

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

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

quantity agreed is Q , then the buyer commits to taking at least $X\%$ of Q every year; if the buyer fails to take delivery of the $X\%$ of Q in a given year, they must nevertheless pay for that LNG. This ensures that the seller's minimum revenue from the contract is the price multiplied with the take-or-pay quantity. A contract may include a destination restriction, which means that the LNG cargo must travel only to the destination that is specified on the contract. Finally, the contract may also include a diversion clause which specifies how the two parties split the surplus in the event that they decide to sell the cargo to a third party (which is known in industry parlance as a "diversion").

Figure 2.1 shows the distribution of contract durations for long-term contracts (defined for the purposes of this study as contracts that are longer than 4 years in duration). The majority of contracts are well over 10 years in length, with the modal contract length being around 20 years.

Long-term contracts do not provide the only way for parties to trade LNG, however. Parties can also trade LNG on the spot market, or using short-term contracts, and this is the focus of this study. A typical spot transaction is negotiated between the buyer and the seller three months prior to delivery date, and will involve the delivery of a single LNG cargo from the seller's export terminal to the buyer's import terminal.³ Figure 2.2 shows that the share of spot and short-term trade has been rising over time, accounting for only 12% of trade in 2004 but 27% of trade in 2017. Finally, buyers can also "re-export" LNG to other buyers by purchasing LNG from one source and re-loading the LNG onto a new tanker. The share of re-exports increased from 0.15% in 2008 to a peak of 2.69% in 2014, but decreased to 1% of total trade by 2017.

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

Figure 2.3 shows the evolution of industry capacity and trade over time. It plots total liquefaction capacity (which is the capacity of LNG exporters), total regasification capacity (which is the capacity of LNG importers) and total LNG trade. The binding constraint on the volume of trade tends to be the available liquefaction capacity, and export capacity utilization tends to be high, ranging from between 80 to 90%. By contrast, there is generally a lot of excess regasification capacity, and this has grown in recent years with increasing entry of new importing countries. These patterns reflect the fact that liquefaction plants are typically more costly to build than regasification plants.


Data

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

Data on liquefaction and regasification capacity I obtain data on investment and capacity data, for both the liquefaction terminals (owned by exporters) and regasification terminals (owned by importers), from the annual reports of the GIIGNL (The International Group of Liquefied Natural Gas Importers), whose members collectively own nearly every LNG import terminal in the world. The dataset covers all liquefaction and regasification terminals operational from 2004 to 2017. The dataset consists of the start-up year, nameplate capacity, the ownership structure and the operator for every export and import terminal built on or prior to 2017. From 2004 onward, I also observe the year when a Final Investment Decision (FID) is made on an export terminal⁴. Due to


⁴Typically the Final Investment Decision (FID) is preceded by various feasibility studies, obtaining the relevant regulatory permits and signing sale and purchase agreements or contracts with buyers. The FID itself is the decision by all project partners to finally commit to the project. Construction of an LNG

Figure 1.1: Histogram of long-term contract durations




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

Figure 1.2: Share of short-term contracts and spot transactions in LNG trade



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

Figure 1.3: Liquefaction capacity, regasification capacity and LNG trade over time



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

time-to-build, the FID year is earlier than the start-up year (which is when the terminal begins operating).

Data on LNG trade flows I obtain data on LNG trade flows from the annual reports of the GIIGNL. This dataset reports yearly LNG trade flows for each country pair from 2004 to 2017. We separately observe trade flows that take place under long-term contracts (exceeding four years in length), short-term contracts and spot trade⁵, and re-exports.⁶ A novel feature of this dataset is that it allows me to distinguish between flows

export terminal only begins once a Final Investment Decision has been taken by the investors.

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

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

that take place under long-term contracts and short-term/spot contracts. Actual trade flows under long-term contracts may differ from the quantity originally agreed to in the contracts between the two countries, either due to production fluctuations or because the parties to the contract agreed to divert cargoes to other destinations. It should also be noted that this dataset record country-country trade flows, rather than firm-firm trade flows.

LNG spot prices and freight costs I obtain data on weekly LNG spot prices and shipping costs between February 2006 and August 2018. The most comprehensive of these datasets comes from Waterborne Energy⁷. The Waterborne Energy dataset reports weekly landed spot LNG prices (measured in USD/MMBtu) at 18 major LNG destinations, as well as weekly freight rates (in USD/MMBtu) for 220 exporter-importer pairs.⁸

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

Data on LNG contracts I also observe data on individual LNG contracts. This

⁷I accessed this dataset through the Reuters Eikon terminal.

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

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

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

⁹The dataset was accessed through the Bloomberg terminal.

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

¹¹An example of a swap is when buyer A originally has a contract with seller A', and buyer B originally

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

Table 3.1 contains summary statistics on key variables used in the analysis. Trade flows and shipping costs are defined at the exporter-importer-year level; export capacity and total exports are defined at the exporter-year level and spot prices and total imports are defined at the importer-year level. A key fact to note is the prevalence of zeros in the trade data: out of 6,406 observations at the exporter-importer-year level, only 1,563 observations feature positive spot trade flows and 1,208 observations feature positive long-run contracted flows. Similarly, out of 2,096 observations at the re-exporter-importer-year, 275 observations feature positive re-exports.¹² The empirical analysis described latter deals with the censoring of the data induced by the presence of zeros.

has a contract with buyer B'. A swap deal would mean that buyer A now purchases from seller B', while buyer B now purchases from seller A'. This kind of arrangement, though, is relatively rare in the industry.

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

Table 1.1: Summary Statistics

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

¹. All trade variables (spot trade, contracted trade, total exports, total imports etc.) are measured in million tonnes or mt. Capacity is measured in million tonnes per annum, or mtpa.


². Spot prices and shipping costs are measured in US\$/MMBtu.

1.3 Descriptive Evidence of Market Power

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


I first discuss a few features of the industry that make it possible in principle for LNG suppliers to exercise market power and price discriminate on the spot market. First of all, the LNG spot market is highly concentrated. Figure 1.4 plots the four-firm concentration ratio over time. It shows that the top four exporters account for over 60% of spot trade in most years, decreasing only from 2014 onward, driven by entry of new LNG suppliers in Australia and the US. The high degree of concentration facilitates the ability of sellers to exercise market power in the spot market.

Figure 1.4: Industry concentration (four-firm concentration ratio) over time



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

Figure 1.5: Shipping costs (\$/MMBtu) for major LNG routes



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

Secondly, shipping costs are important in this industry. High shipping costs cause markets to become segmented from one another and make it feasible in theory for sellers to price discriminate across markets. As Figure 1.5 illustrates, the freight rate varies considerably over time, driven by both changes in the crude oil price (which determines the fuel cost incurred in LNG shipping) and changes in LNG demand: during periods of high LNG demand, demand for LNG shipping increases and this leads to higher freight rates. When shipping costs are high, sellers face less competition both from faraway sellers and faraway buyers who can engage in arbitrage behavior. For example, as Figure 1.5 shows, shipping costs were at historically high levels between 2012 and 2014. During this period, Qatar (the largest LNG producer) faced less competition when selling to Asia (the high-demand market during this period), as some of its major competitors (Nigeria and Trinidad) incurred much higher transportation costs when selling into Asia. Moreover, the high transportation costs increased arbitrage costs for European buyers looking to re-export LNG from Europe to Asia. This made it more feasible for Qatar as well as other sellers located closer to Asia to charge higher prices when selling to Asian buyers.

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

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

already selling to the third destination has little incentive to allow such cargo diversions, since the result would be that the seller faces competition in that market from their own contracted buyer.

LNG buyers do have the ability to receive a contracted LNG cargo at their liquefaction terminal, reload the LNG onto a new tanker and then re-sell the cargo to a new destination. The costs of engaging in such re-exports can, however, be high. They include not just the shipping cost between the original and the new destination, but also a fee paid to the terminal operator as well as a host of logistical constraints that further add to the costs of re-exporting.¹⁴ As a result, an industry report on re-exports argues that “These logistical factors mean that the true cost of reload is well above the direct cost paid to the terminal operator. In other words a premium well in excess of shipping costs is required to incentivize the re-exporting of gas.”¹⁵

The above discussion suggests that the *possibility* for sellers to exercise market power exists in this industry. At the same time, it is not a priori obvious just from these facts that sellers would exercise market power. The firms involved in LNG exporting typically have close ties to national governments: indeed, in the majority of LNG exporting countries, LNG exporting companies are majority-owned by national oil and gas companies.¹⁶ Thus it is not immediately apparent that LNG exporters’ behavior on the spot market would conform to short-run profit maximization as in standard oligopolistic models. One could imagine too that sellers might have motives to refrain from price discriminating as a way to build relations with buyers, given that most of

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

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


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

their sales still come from long-term contracts. Finally, though the LNG industry is significantly less concentrated on the buyer side than on the seller side, there do exist large buyers, unlike in standard settings of oligopolistic behavior. As such, whether firms exercise market power and price discriminate is ultimately an empirical question that needs to be answered by looking at how firms actually behave on the spot market, rather than assumed from the outset.

I now describe the recent evolution of spot prices and spot trade flows in the industry, and argue that these patterns are difficult to reconcile with competitive behavior on the part of the sellers. Figure 2.6 shows spot prices in various LNG importing regions. The feature that stands out from this figure is the presence of significant and systematic spot price differentials across regions, especially during periods when the LNG market is tight. A particularly striking illustration of this is the period between mid-2011 and end-2013, when there was a large spike in Japan's LNG demand following the Fukushima nuclear disaster. Asian spot prices, as well as spot prices in Latin America, remained an average of \$5/MMBtu higher than European prices during this period. The prices converge again only in late 2014 and early 2015, driven by a combination of reduced LNG demand in Asia, lower transportation costs (driven by the oil price collapse) and addition of new export capacity starting from 2015. In addition to the post-Fukushima period from 2011 to 2013, there have been other periods when prices in different regions sharply diverged from one another. For example, between January 2007 and July 2008, another period of tight demand, Asian spot prices were about \$3.5/MMBtu higher on average than European spot prices.


In a competitive spot market with no frictions, faced with these divergent prices, sellers would sell only to the destination with the highest price net of transportation costs. This would suggest that during these periods of high Asia-Europe spot price

Figure 1.6: LNG Spot Prices in Different Regions



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

Figure 1.7: Short-run and spot LNG imports to different regions



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

differentials, spot exports should be mostly directed to Asia, as even after accounting for transportation costs most sellers received higher prices from selling to Asia than to Europe. As Figure 2.8 shows, though, during the “boom” period of 2011-14 Europe continued to import a significant amount of LNG on the spot market.¹⁷ Moreover, much of Europe’s spot market imports during this period came from sellers that would have received higher prices from selling to Asia. For example, in 2012, 46% of Europe’s spot market purchases were from Qatar, and a further 12% were from Egypt and Algeria. This is in spite of the fact that these countries faced roughly similar shipping costs of selling LNG to Asia and Europe (see Figure 1.5 for shipping costs from Qatar to Asia and Europe), meaning the marginal cargo garnered a premium of \$2-4/MMBtu if it were sent to Asia rather than Europe. The fact that sellers in Middle East and North Africa continued to sell to Europe is inconsistent with competitive behavior on the part of LNG exporters.

The pattern of sellers selling to multiple markets at different received prices suggests that sellers deviate from perfectly competitive behavior. However this still does not establish whether these deviations are the result of market power or some other sources, such as unobserved trade frictions. I now present a reduced-form test for whether sellers exercise market power. The test is designed to distinguish the role of market power from other trade frictions. The intuition behind the test is that if sellers exercise market power, small and large sellers should allocate their sales differently across markets. LNG sellers have the choice of selling to several markets, each of which has different demand shocks. If the seller is small in size, it should behave almost competitively and sell most of its output to the market from which it receives the highest price (net

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

of shipping costs). If the seller is large and exercises market power, though, it has an incentive to withhold LNG from the market with high demand and sell more LNG to the market with low demand, creating a wedge in prices between the two destinations. Thus, if sellers exercise market power, we should see large sellers on the spot market respond less strongly to price differentials than small sellers. Under the alternative hypothesis of perfectly competitive behavior, even with trade frictions, both small and large sellers should be equally responsive to spot price differentials.

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

$$\ln(S_{ijt}/S_{it}) = \alpha + \beta \ln(p_{ijt}) + \gamma \ln(p_{ijt}) * \ln(S_{it}) + X_{ijt} * \theta + \eta_{ijt} \quad (1.1)$$

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

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

Table 1.2: Reduced Form Evidence for Market Power: Regression of Share of Spot Sales to market j on Price at market j

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

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


Table 1.2 shows the estimated coefficients from this regression. As expected, the coefficient on the logarithm of prices is positive: holding all other prices fixed, if a demand shock causes the price at destination j increase by 10%, the share of spot output sold to that market increases by 4.4%. However the coefficient on the interaction term between price and size is negative, implying that the larger the total spot sales of seller i , the lower is seller i 's quantity response to a price increase in destination j . Moreover, comparing columns (1) and (2), there is a noticeable increase in the R^2 when we include this term, suggesting that a large part of the variation in spot sales is driven by market

power.

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

Finally, some additional evidence suggestive of seller market power is provided by Figure 1.8, which plots the average number of buyers each seller sells to on the spot market against its total spot market sales. There is a positive correlation between seller size and the number of buyers to which the seller sells, suggesting that large sellers are more likely to price discriminate by splitting their sales across multiple buyers, while small sellers generally sell to a small number of buyers that offer the highest price.

Figure 1.8: Relationship between total spot sales and number of countries that a seller exports to



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

1.4 Model

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

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

quantities. Under the hypothesis of perfect competition, sellers choose their spot quantities taking equilibrium prices as given. The model nests both of these models of seller behavior.

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

1.4.1 Model Description

Buyers

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

$$Q_j = Q_j^d(p_j, \varepsilon_j) \quad (1.2)$$

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

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

$$Q_j^d(p_j, \varepsilon_j) = \varepsilon_j - bp_j \quad (1.3)$$

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

Sellers

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

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

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

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

Costs of production and sales

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

Let $C(q_i, k_i)$ denote the cost of production, which I will assume is continuously differentiable in both arguments. The inclusion of k_i in the cost function reflects the

fact that capacity constraints are important in this industry (as evidence by Figure 2.3) and sellers with larger capacity are able to export more LNG.

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

Thus firm i 's cost function is:

$$C_i(q_{i1}, \dots, q_{ij}, \dots, q_{iJ}, k_i) = \underbrace{C(q_i, k_i)}_{\text{Cost of production}} + \underbrace{\sum_j (d_{ij} + \phi(x_{ij}) + v_{ij}) q_{ij}}_{\text{Cost of trading}} \quad (1.4)$$

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

Re-sales

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

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

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

Information structure

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

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

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

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

Seller competition

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

I consider two possible models of seller behavior. The first model allows sellers to

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

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

1.4.2 Equilibrium

Market clearing

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

$$Q^d(p_j^*, \varepsilon_j) = \sum_{i=1}^N q_{ij}^c + \sum_{i=1}^N S_{ij}(v_{ij}), \forall j \quad (1.5)$$

Optimal behavior

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

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

$$\max_{\hat{S}_{ij}} E \left[\sum_{k=1}^J p_{ik}^c q_{ik}^c + \sum_{k \neq j} p_k^* S_{ik}(v_{ik}) + p_j^* \hat{S}_{ij}(v_{ij}) - C(q_i, k_i) - \sum_k (d_{ik} + \phi(x_{ik}) + v_{ik}) q_{ik} \right] \quad (1.6)$$

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

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

$$\begin{aligned} & \max_{\hat{S}_{ij}} E \left[p_j^* \hat{S}_{ij}(v_{ij}) - C(q_i, k_i) - (d_{ij} + \phi(x_{ij}) + v_{ij}) \hat{S}_{ij}(v_{ij}) \right] \\ & \max_{\hat{S}_{ij}} E (p_j^* | \hat{S}_{ij}(v_{ij})) \hat{S}_{ij}(v_{ij}) - E(C(q_i, k_i) | \hat{S}_{ij}(v_{ij})) - (d_{ij} + \phi(x_{ij}) + v_{ij}) \hat{S}_{ij}(v_{ij}) \quad (1.7) \end{aligned}$$

where equation (1.7) makes it clear that the seller faces uncertainty both over the price in market j as well as its total costs (since it does not know exactly how much it will sell in other markets).

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

$$E(p_j^*|S_{ij}(v_{ij})) + S_{ij}(v_{ij}) \frac{\partial E(p_j^*|S_{ij}(v_{ij}))}{\partial S_{ij}} - \frac{\partial E(C(q_i, k_i)|S_{ij}(v_{ij}))}{\partial S_{ij}} - (d_{ij} + \phi(x_{ij}) + v_{ij}) \leq 0 \quad (1.8)$$

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

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

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

$$E(p_j^*|S_{ij}(v_{ij})) - \frac{\partial E(C(q_i, k_i)|S_{ij}(v_{ij}))}{\partial S_{ij}} - (d_{ij} + \phi(x_{ij}) + v_{ij}) \leq 0 \quad (1.9)$$

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

Both cases can be compactly summarized by introducing a conduct parameter $\theta_i^c \in \{0, 1\}$ (Bresnahan, 1982), as in equation (1.10). When $\theta_i^c = 0$, firms behave competitively. When $\theta_i^c = 1$, firms engage in Cournot competition:

$$E(p_j^*|S_{ij}(v_{ij})) + \theta_i^c S_{ij}(v_{ij}) \frac{\partial E(p_j^*|S_{ij}(v_{ij}))}{\partial S_{ij}} - \frac{\partial E(C(q_i, k_i)|S_{ij}(v_{ij}))}{\partial S_{ij}} - (d_{ij} + \phi(x_{ij}) + v_{ij}) \leq 0 \quad (1.10)$$

A key goal of the estimation exercise will be to estimate equation (1.10) to see if the actual behavior of sellers is better characterized by perfect competition or by Cournot competition.

Equilibrium re-sales

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

$$\forall j, j', q_{j'j}^s = \max\left(0, \frac{1}{\delta_r}(p_j - p_{j'} - d_{j'j} - \kappa)\right) \text{ if } \sum_j q_{j'j}^s < Q_{j'} \quad (1.11)$$

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

In Appendix A.1, I describe Monte Carlo simulations of a simplified version of the above model that I carried out in order to illustrate the effect of market power in this model. In the simplified version of the model, trade costs are deterministic and marginal costs are linear until production hits a binding capacity constraint. I first consider a scenario where firms have large amounts of excess capacity. I show that market power results in both lower overall production and greater splitting of sales across buyers, with sellers more likely to sell to faraway buyers than under competitive behavior. I then consider a scenario where firms are highly capacity-constrained, so that they cannot adjust their total production. I show that in this case, market power still leads to lower welfare because the allocation of quantities across markets is inefficient. Sellers sell

more of their output to faraway destinations under Cournot competition. When there is a demand shock in one market, sellers do not increase their sales to that market as sharply as under competitive behavior, which results in higher prices in the high-demand market than in the low-demand market and misallocation across the two markets.

1.5 Estimation

I now describe how to estimate the model described in Section 2.4. I allow the data to identify whether sellers behave strategically as in Cournot competition or non-strategically as in perfect competition. I do so by estimating a conduct parameter that nests various models of seller behavior, together with parameters describing seller costs. A precondition for estimating cost and conduct parameters is knowledge of the industry demand curve. As such I begin by describing how I estimate the demand curve, and then follow up by describing the method for estimating costs together with the conduct parameter. For now I abstract away from estimating buyer re-sale costs, but plan to incorporate resale cost estimation into future versions of the paper.

1.5.1 Estimation of Demand Curve

I estimate the following demand curve for importing country j :

$$Q_{jt} = Q_j^d(p_{jt}, X_{jt}, \varepsilon_{jt}) = \alpha - bp_{jt} + X_{jt}\gamma + \varepsilon_{jt} \quad (1.12)$$

Here X_{jt} denotes demand shifters. In the empirical analysis, this includes electricity consumption as well as country fixed effects. Liquefied natural gas is almost always converted back into gaseous form before it is consumed, and natural gas is primarily

used for power generation and heating. Electricity consumption thus captures changes in electricity demand that translate into a shift in the demand curve for LNG. I use country fixed effects to capture differences in heating demand across countries as well as other time-invariant country-specific factors that affect LNG demand (for example the amount of piped natural gas the country has access to).

Table 1.3: Demand Curve Estimates

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

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

Estimation of equation (2.28) via ordinary-least squares runs into standard endogeneity problems because the spot price p_{jt} is positively correlated with the unobserved demand shock ε_{jt} : the greater country j 's demand for LNG, the higher the equilibrium spot price in country j . I consider two instruments for p_{jt} . The first instrument, which varies only over time but not across j , is total liquefaction capacity (i.e. total export capacity) in the world in period t . The greater LNG capacity in period t , the higher is the supply of LNG and therefore the lower the price in period t . The identification assumption is that LNG export capacity is uncorrelated with idiosyncratic

demand shocks today, after controlling for electricity consumption. The logic behind the instrument is that LNG terminals take many years to build, and at the time the decision to invest is made, it is difficult to foresee idiosyncratic demand shocks that are realized several years later. The modern history of the LNG industry is replete with instances where sellers make investments without fully anticipating how demand would evolve in the importing countries. For example, Qatar's massive capacity expansion in the 2000s was driven to a large degree by the expectation that the US would be a major importer of natural gas. However by the time all of Qatar's terminals came online, US demand for LNG had shrunk dramatically due to the shale gas boom, and instead Qatar ended up turning to Asian countries as its major buyers of LNG.

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

The demand curve estimates are shown in Table 2.4. Column (1) shows the OLS regression of demand on spot prices, without any controls. As expected, the coefficient is positive, reflecting the fact that price changes are mostly driven by shifts in demand rather than shifts in supply. In Column (2)-(4), I instrument for spot prices using total liquefaction capacity and total electricity demand in countries other than country j . The 1st-stage Cragg-Donald Wald F-statistic ranges from 19 to 35, suggesting that weak instruments are unlikely to be a concern. In the most flexible specification (4) where I control for both electricity demand and importer fixed effects, the coefficient on the

spot price is negative and statistically significant, and implies a demand elasticity of around 0.33 for the average observation. The coefficient on electricity consumption is also positive and statistically different from zero, and implies that the elasticity of LNG demand with respect to electricity consumption is around 0.82 for the average observation.

1.5.2 Estimation of Cost and Conduct Parameters

The starting point for the estimation of the cost and conduct parameters is the first-order condition (1.10), which I have repeated here:

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

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

$$C(q_i, k_i) = c_i q_i + \frac{\delta(k_i)}{2} q_i^2 \quad (1.13)$$

I show in Appendix A.2 that under the assumption of linear demand and quadratic costs, the first-order condition can be written as:

$$\begin{aligned} \beta(\varepsilon_j - Q_j^c - \sum_{i' \neq i} E(S_{i'j}(v_{i'j}))) - (c_i + \delta(k_i)(q_i^c + \sum_{k \neq j} E(S_{ik}(v_{ik}))) + d_{ij} + \phi(x_{ij})) - \\ (\beta(1 + \theta_i^c) + \delta(k_i)) S_{ij}(v_{ij}) - v_{ij} \leq 0 \end{aligned} \quad (1.14)$$

Let A_{ij} be defined as follows:

$$A_{ij} = \beta(\varepsilon_j - Q_j^c - \sum_{i' \neq i} E(S_{i'j}(\mathbf{v}_{i'j}))) - (c_i + \delta(k_i)(q_i^c + \sum_{k \neq j} E(S_{ik}(\mathbf{v}_{ik}))) + d_{ij} + \phi(x_{ij})) \quad (1.15)$$

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

The FOC can be written as:

$$A_{ij} - (\beta(1 + \theta_i^c) + \delta(k_i))S_{ij}(\mathbf{v}_{ij}) - \mathbf{v}_{ij} \leq 0 \quad (1.16)$$

Define

$$S_{ij}^* = \frac{1}{\beta(1 + \theta_i^c) + \delta(k_i)}(A_{ij} - \mathbf{v}_{ij}) \quad (1.17)$$

Then the firm's policy function is:

$$S_{ij} = \max(S_{ij}^*, 0) \quad (1.18)$$

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

The firm's expected sales, before observing $\mathbf{v}_{ij}$, are given by:

$$\begin{aligned} E(S_{ij}(\mathbf{v}_{ij})) &= Pr(\mathbf{v}_{ij} < A_{ij})E(S_{ij}^* | \mathbf{v}_{ij} < A_{ij}) \\ &= \frac{1}{\beta(1 + \theta_i^c) + \delta(k_i)} Pr(\mathbf{v}_{ij} < A_{ij})(A_{ij} - E(\mathbf{v}_{ij} | \mathbf{v}_{ij} < A_{ij})) \end{aligned} \quad (1.19)$$

Equation (1.19) shows how market power affects seller behavior. When the conduct parameter is 0, the slope of the firm's sales with respect to A_{ij} is large, meaning that firms are very responsive to price differentials across regions. By contrast when the conduct parameter is 1, the slope is smaller, meaning that firm sales do not adjust perfectly

to pricing differentials, and firms price discriminate in equilibrium by charging higher prices to buyers to whom they have higher spot sales. Thus the elasticity of spot sales with respect to the net price (price minus shipping cost) is what identifies the Cournot model separately from the competitive model. I have used the model to generate many simulated datasets of the same sample size as the dataset I use for estimation, and found that in general I can estimate the conduct parameter quite precisely.

Estimation routine

The parameter vector I have to estimate is $\theta = (\sigma, \phi(\cdot), c_i, \delta(K), \theta_i^c)$. I estimate the model via a maximum likelihood procedure. For each guess of θ , I solve for the policy thresholds A_{ij} , and then use equations (1.17) and (1.18) to form the likelihood. I search for the θ that leads to the highest sample likelihood function. Notice that this is essentially a Tobit model where the censoring of spot quantities at zero is explicitly modeled.

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

- Begin with an initial guess of $A_{ij} \forall i, j$.
- At step $(k + 1)$, we know the policy threshold A_{ij}^k from the previous step k . Then for each i, j :
 - Compute $\sum_{i' \neq i} E(S_{i'j}(v_{i'j}))$ and $\sum_{k \neq j} E(S_{ik}(v_{ik}))$ using A_{ij}^k and equation (1.19).
 - Update the policy threshold to A_{ij}^{k+1} using equation (1.15).

- Continue iterating till $||A_{ij}^{k+1} - A_{ij}^k|| < \varepsilon$, where ε is a sufficiently small threshold.

Results

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

In the main specification (column (2) of Table 2.6), each buyer is a separate importing country. In an alternative specification, I aggregate buyers into nine importing regions: North East Asia, Southeast Asia, South Asia, Middle East, Southern Europe, Northern Europe, Central America, North America and South America.

Table 2.6 shows that the estimated conduct parameter is estimated to be 0.72 (when we treat each importing country as a separate market) and 0.90 (when we treat each importing region as a separate market). A conduct parameter of 1 would indicate that firms are behaving exactly as the Cournot model would predict. The conduct parameter estimate of 0.72-0.90 suggests that the Cournot model may be a reasonable description of how sellers choose spot quantities in practice, though it also suggests that seller behavior may be more competitive than Cournot behavior. It appears unlikely that sellers engage in competitive behavior, since the estimated conduct parameter is very far from 0.

The estimated δ parameter suggests that increasing capacity utilization by 10% increases marginal costs by \$1.5/MMBtu, compared to an average price of \$8.6/MMBtu.

This indicates that capacity constraints are an important component of marginal costs in this industry.

Table 1.4: Conduct and Cost Parameter Estimates		
	(1)	(2)
SD of match shocks, σ	3.08	3.72
Conduct parameter, θ^c	0.90	0.72
Convexity, δ	7.47	8.59
N	1342	3353
Regions	9	40
Fit (R^2)		
Prices p_j	0.66	0.70
Total Exports by Exporter q_i	0.98	0.98
Total Spot Exports by Exporter $\sum_j S_{ij}$	0.73	0.74
Spot Trade Flows S_{ij}	0.55	0.46

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


Model Fit

The model fits many of the key data patterns reasonably well, as shown in the last panel of Table 2.6. The R^2 for prices is 0.66-0.70 (depending on the exact specification), while the R^2 for spot trade flows is 0.46-0.55. The R^2 for total exports is close to 1, but this is because total exports are largely pinned down by contracts. Even so, when we look at total spot exports by seller, the R^2 is reasonably high at around 0.73-0.74.

Figure 2.12 shows actual versus predicted prices in Northeast Asia and South-western Europe, using specification (2) in Table 2.6. Though the model does not predict prices precisely, it is able to predict both the small price differentials prior to 2011,

Figure 1.9: Actual vs. Predicted Spot Prices

Actual vs. Simulated Prices, Regional Averages



Note: The figure plots actual and predicted annual prices in Northeast Asia (comprising Japan, China, Korea and Taiwan) and South-western Europe (comprising Spain, France, and Italy). The predicted prices are calculated by simulating the estimated model (specification (2) in Table 2.6) for 1000 different match shocks, and averaging across each of these simulations.

and the fact that prices in Asia diverged from prices in Europe starting from 2011 and lasting till 2014. One feature of the data that the model does less well at predicting is the extent of the price convergence from 2015 onward: though the model-predicted price differentials do shrink from 2015 onward, they still remain substantially higher compared to the true price differentials. This period coincided with the building of new capacity in Australia and the US as well as a drop in shipping costs. There are two possible reasons why the model does less well in the last three years of the sample. One is that it treats exporting countries as decision-making entities. This might be a reasonable assumption for countries such as Qatar, Indonesia and Malaysia (where LNG exporting is under the control of a national oil company), but likely to be less tenable for Australia and US (where each terminal operates as a separate decision-making entity), and these are precisely the two countries that expand capacity in the latter part of the sample. A second possibility is that perhaps the conduct parameter is not fixed over

time, and the market became more competitive after the entry of new suppliers, so that suppliers acted more competitively than the Cournot model would predict after 2014.


Consequences of market power

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

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

Figure 1.10 shows price differentials across regions in the baseline simulation with market power and the counter-factual simulation without market. We find that price differentials are much smaller when we remove seller ability to price discriminate. This effect is especially pronounced in the period between 2011 and 2013, following the Fukushima disaster in Japan. We find that the price differential between Asia and Europe would average around \$2.5 between 2011 and 2013 if sellers behaved competitively, versus a price differential of around \$4.5 when sellers price discriminate. Thus 45% of the post-Fukushima price differential between Asia and Europe is explained by seller market power in the LNG spot market.

Figure 1.10: Price Differential between Northeast Asia and Southwestern Europe: Baseline vs. Competitive



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

1.6 Conclusion

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

seller decisions are consistent with a Cournot model and unlikely to be generated by a competitive model.

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

CHAPTER 2

**LONG-TERM CONTRACTS AND EFFICIENCY IN THE LIQUEFIED
NATURAL GAS INDUSTRY**

Nahim Bin Zahur

Abstract: Long-term contracts facilitate efficient levels of investment by sellers faced with the risk of ex-post holdup. Contractual rigidities, however, reduce the ability of the market to respond flexibly to demand uncertainty. This paper provides the first empirical analysis of the trade-off between hold-up risk and contract rigidity, focusing on the liquefied natural gas (LNG) industry, where long-term contracts account for over 70% of trade. I structurally estimate a model of contracting, investment and spot trade that incorporates hold-up risk and contractual rigidities, using a rich dataset on LNG contracts, investment, trade flows, and spot prices. LNG buyers are estimated to have substantial bargaining power, implying that sellers face considerable hold-up risk when making investment decisions. In the short-run, allocative efficiency would improve through reduced use of long-term contracts. However, investment decreases by 35% in the absence of long-run contracting, suggesting that long-term contracts play a significant role in combatting holdup at the cost of short-term allocative efficiency.

2.1 Introduction

Long-term contracts are widely used to carry out business-to-business transactions in capital-intensive industries. A classical rationale behind long-term contracts is that they facilitate sunk cost investments (Williamson, 1975, 1979; Klein *et al.*, 1978; Grossman and Hart, 1986). If agents invest before negotiating the terms of trade, they may

be unable to recoup the full cost of the investment in ex-post bargaining. Signing a long-term contract before investing can protect the investor from the risk of ex-post holdup. However, a drawback of long-term contracts is that they may be inflexible in response to fluctuations in demand and costs, as it is costly and difficult to account for all possible contingencies when writing a contract (Masten and Crocker, 1985). By contrast, using the spot market to allocate goods allows firms greater flexibility in responding to uncertainty, which in turn can lead to greater allocative efficiency.

The trade-off between hold-up risk and contract inflexibility is particularly salient in the global liquefied natural gas (LNG) industry. LNG constitutes one of the fastest-growing energy markets in the world, with the size of the market doubling between 2004 and 2017. An increasing number of countries have turned to LNG imports as a way to diversify their energy sources and reduce their reliance on piped imports of natural gas. The volume of LNG trade is largely determined by the available capacity of liquefaction terminals that convert natural gas into LNG. Investment in liquefaction terminals thus has crucial implications for prices and welfare. Constructing a terminal, however, requires sellers to incur large upfront costs typically in excess of \$10 bn. Under-investment due to hold-up is therefore a natural concern, and sellers have historically only been willing to make these investments after signing long-term contracts with buyers. At the same time, the LNG market has been characterized by demand fluctuations, such as the Fukushima nuclear crisis in 2011 which led to a sharp increase in Japan's LNG demand and a 50% increase in Asian LNG prices. Long-term contracts are potentially inflexible in dealing with such demand shocks, especially because contracts often include clauses that limit resales.

Quantifying the trade-off between hold-up risk and contract rigidity is important in order to assess the benefits and costs of long-term contracting. However, this is

a challenging problem: detailed data on long-term contracts is often lacking, and the existing literature does not provide a tractable empirical framework for analyzing long-term contracting. In this paper, I fill this gap by developing an empirical model of LNG contracting, investment and spot trade that features both hold-up risk and contract rigidity. I propose a novel estimation strategy that leverages the timing of contracting and investment decisions to distinguish between hold-up and other motives for contracting (such as energy security concerns). My approach allows the researcher to learn about bargaining power in the absence of data on contract prices, a common data challenge with studying long-term contracts. Using the estimated model, I study the consequences of long-term LNG contracts for investment and allocative efficiency.

The model that I develop features a set of spatially differentiated sellers and buyers. Buyers have stochastic demand for LNG which cannot be perfectly predicted in advance. I allow for buyers to have a different willingness-to-pay for contracted and spot LNG, as a way to capture energy security motives that buyers often cite as a reason for signing long-term contracts. Sellers invest in LNG terminals that allows natural gas to be converted into LNG. Once a terminal has been built, sellers export LNG to buyers, incurring production costs that are a function of their capacity and shipping costs that depend on the distance between the buyer and seller.

Trade takes place via two distinct mechanisms. A seller and a buyer can sign a long-term contract which specifies a fixed quantity to be traded between the buyer and the seller each year, as well as a price paid by the buyer to the seller. I model the contracting decision as a simultaneous Nash bargaining game played between the seller and buyers, where they bargain over both quantities and prices to be paid by the buyer to the seller, building on Chippy and Snyder (1999). Alternatively, the sellers and buyers may trade repeatedly via a spot market that takes place every year. In the spot market, sellers

observe demand shocks and engage in Cournot competition when choosing their spot sales to each region, taking as given all contracted quantities. As in Miller and Osborne (2014), sellers engage in spatial price discrimination, charging different spot prices to buyers in different regions.

The model captures the key trade-off between hold-up risk and contract rigidity that firms in the industry face when deciding whether to trade using long-term contracts or the spot market. Sellers face hold-up risk because the investment cost has to be incurred upfront and they may not be able to recover a return that justifies the cost of investment. As such, the seller has an incentive to sign long-term contracts with buyers before they have committed to the investment. The incentive is stronger the weaker the seller's bargaining position relative to the buyer; for example, when the seller is located far away from alternative buyers (so that the seller is effectively forced to transact with one buyer). At the same time, contractual inflexibility can be costly because long-term contracts lock in transactions before the parties have full information about demand. This restricts the ability of sellers to meet demand shocks by reallocating LNG across buyers. When demand uncertainty is high enough, the use of long-term contracts can result in a reduction in allocative efficiency.

I carry out the analysis using a rich dataset on capacity, long-term contracts, trade flows, and spot prices in the LNG industry that I have collected from a variety of industry sources. The dataset has several unique features that I leverage in the empirical analysis. Firstly, I observe the precise timing of contract signature decisions and investment decisions, which allows me to classify contracts into those that are signed prior to the date of investment, and those that are signed after the date of investment. The distinction between pre-investment and post-investment contracting is important for learning about hold-up, since the risk of holdup only arises if bargaining occurs after the seller has

committed to the investment. Secondly, in addition to data on long-term contracts and capacity, I also observe rich information on the LNG spot market, including spot LNG trade flows, spot prices and shipping costs. This data is particularly useful for empirically modelling the LNG spot market.

A number of reduced-form data patterns suggest both that the hold-up effect drives contracting behavior, and that contractual rigidities contribute to short-run misallocation. First, much of contracting activity takes place before investment: on average, a seller signs contracts amounting to around 60% of their capacity before making a final investment decision. Moreover, the further away sellers are located from buyers, the greater the contract quantity they sign before committing to investment. This is consistent with a hold-up story, where sellers that are located far away from potential buyers have lower bargaining leverage and compensate by signing larger contracts prior to investing. Secondly, the LNG market in recent years has been characterized by large and systematic price differentials across regions that are difficult to rationalize based on competitive behavior. I find that sellers with a larger share of their capacity committed to long-term contracts are less responsive to short-run price differentials, suggesting that contract rigidities contribute to short-run misallocation.

I then proceed to estimate the structural model. I begin by estimating demand and short-run production costs. I first estimate demand curves for each buyer, using aggregate export capacity and electricity demand in other regions as instruments for the spot price. I then use data on LNG spot trade flows, spot prices and shipping costs to estimate seller production costs via a minimum distance procedure, where I compute the model's predicted vector of spot trade flows for each parameter guess and identify the parameter guess that brings the predicted trade flows as close as possible to the actual spot trade flows in the data. My estimates of buyer demand and production costs

suggest that short-run demand is both inelastic and uncertain and that sellers are highly capacity-constrained in the short-run. The estimated model is able to fit the large inter-regional price differentials that are seen in the data and that are indicative of short-run misallocation.

Next I estimate a set of structural parameters characterizing the contracting and investment decisions, including the cost of investment, contract preference parameters and a bargaining weight parameter that governs the distribution of surplus between seller and buyer. I show that the logic behind the hold-up effect can be inverted in order to identify the bargaining power of sellers relative to buyers: since sellers with low bargaining power have a stronger incentive to engage in a greater degree of pre-investment contracting, the extent to which sellers sign contracts prior to investment helps identify their relative bargaining weight. This allows me to learn about bargaining power despite not observing the prices of LNG traded under long-term contracts.

I find that bargaining power is split close to evenly between the buyers and sellers, and am able to reject the hypothesis that sellers have the ability to make take-it-or-leave-it offers. Sellers face a substantial degree of hold-up risk when investing: on average, the marginal value of increasing the capacity of an LNG project would be 20% higher if sellers were able to fully extract the surplus from investing. I also find that buyers value contracted LNG by around \$0.7/MMBtu (or roughly 8% of average spot prices) more than spot LNG. Finally investment costs are estimated to be concave indicating there are economies of scale in project construction.

Using the estimated model, I carry out two counter-factual exercises in order to explore the consequences of long-term contract for efficiency. In the first counter-factual, I hold industry capacity fixed and reduce the quantity traded under every long-term contract by 50%. This allows me to measure the allocative efficiency consequences

of using long-term contracts. I find that over the period from 2006 to 2017, aggregate welfare would increase by \$27 bn relative to the benchmark scenario, or around 2% of industry revenue. The gains from increased flexibility would be more than twice as large (at around \$70 bn) if the spot market were competitive: as pointed out by Allaz and Vila (1993), contracts have a pro-competitive effect when sellers have market power on the spot market, and so reducing the degree of contracting leads to sellers having a stronger incentive to price discriminate on the spot market. However I find that the net benefits from greater flexibility outweigh the potential losses from increased market power.

In the second counter-factual, I shut down the ability of sellers and buyers to sign long-term contracts, and solve for each sellers' investment decisions (where they can only sell LNG on the spot market). I find that the average seller reduces investment by 35%. The largest reduction in investment in this counter-factual comes from sellers who are on average located further away from buyers than the median seller: these are the sellers that are most subject to hold-up risk, and as such are the ones that are least likely to invest without the protection afforded by long-term contracts.

This paper contributes to three main strands of literature. First, it builds on an extensive literature on the theoretical foundations of transaction-cost theory (Williamson, 1975, 1979; Klein *et al.*, 1978; Grossman and Hart, 1986) and empirical tests of the theory (Joskow, 1985, 1987; Crocker and Masten, 1988; Hubbard, 2000; Brickley *et al.*, 2006). To my knowledge, this is the first paper to provide a empirical framework for analyzing how firms' contracting choices are influenced by hold-up risk, and quantifying the effect of contracting choices on investment. An additional contribution I make to this literature is that I use the timing of contracting and investment decisions to learn about the hold-up effect; by contrast, the existing empirical literature generally uses the duration of contracts to test predictions of transactions-cost theory.¹

¹For example, Joskow's seminal studies of the US coal market (Joskow, 1985, 1987) find, consistent with

The paper contributes to a growing literature in industrial organization that uses bargaining models of industry oligopolies to study negotiations between firms (Horn and Wolinsky, 1988; Chipty and Snyder, 1999; Crawford and Yurukoglu, 2012; Grennan, 2013; Gowrisankaran *et al.*, 2015; Ho and Lee, 2017; Goetz, 2019). Most of the papers in this literature infer bargaining power either from negotiated prices (e.g. Crawford and Yurukoglu (2012); Grennan (2013); Gowrisankaran *et al.* (2015); Ho and Lee (2017)), or from bargaining delays (e.g. Goetz (2019)). My paper provides a new strategy for inferring bargaining power from the timing of contracts and investment, using the insight that firms with lower bargaining leverage are more exposed to hold-up risk when investing and therefore have an incentive to sign larger contracts before investing. The strategy can be useful in other settings where the analyst does not observe negotiated prices but does observe investment and contracting decisions that are functions of these negotiated prices.

Finally, my paper contributes to the literature measuring the effects of contracting on short-run allocative efficiency in energy markets. Building on the theoretical insights from Allaz and Vila (1993), a series of papers have empirically analyzed electricity markets in the U.S. and elsewhere and have found that the use of forward contracts significantly mitigates the deadweight loss from seller market power (Wolak, 2003; Bushnell *et al.*, 2008; Ito and Reguant, 2016). In contrast to this literature, I find that allocative efficiency can decrease from the use of contracts, despite the pro-competitive effect of contracts (which also exists in my setting), due to the inflexibility of long-term contracts in responding to demand uncertainty. In addition, the paper is also related to a growing literature studying long-term contracting in the LNG industry (Brito

transactions-cost theory, that “minemouth” coal projects which are located right next to the coal mine sign contracts of significantly longer duration than other projects. Crocker and Masten (1988) find in the US natural gas industry that firms signed longer contracts when there was a greater likelihood of hold-up but signed shorter contracts when there was greater uncertainty.

and Hartley, 2007; Neumann, 2009; Ruester, 2009; Hartley, 2015; Agerton, 2017). In particular, Brito and Hartley (2007) and Hartley (2015) develop theoretical models to examine the trade-offs faced by parties when deciding whether to sign long-term contracts or trade on the spot market. My paper builds on this literature by estimating a new structural model of LNG contracting, investment and trade in order to empirically quantify these trade-offs.²

The remainder of the paper is divided as follows. Section 3.2 discusses key institutional features of the LNG industry and describes the dataset I use for the analysis. Section 2.3 discusses descriptive evidence, including evidence for contracting rigidities. Section 2.4 develops a model of contracting, investment, trade and pricing in this industry. Section 2.5 describes estimation of the key parameters of the model. Section 3.6 discusses counter-factual exercises that I carry out using the estimated model. Section 2.7 concludes.

2.2 Industry and Data

Liquefied natural gas (LNG) is natural gas that has been cooled into a liquid form so that it can be transported in specialized LNG tankers. For LNG trade to take place, both the exporting country and the importing country need to invest in specialized infrastructure. In the exporting country, a liquefaction terminal converts natural gas into LNG, and loads the LNG onto a tanker. The tanker then transports LNG to a receiving terminal (known as a regasification terminal) at the importing country, where the LNG

²It should also be noted that Hartley (2015) emphasizes the role of long-term contracts in increasing the debt capacity of investment projects by reducing cash flow variability. By contrast, this paper follows in the tradition of the transaction-cost literature (such as Williamson, 1975) and focuses on quantifying how long-term contracts reduce the hold-up risk faced by agents making sunk cost investments.

is unloaded. Once the LNG has been unloaded, it is usually converted back into gaseous form and used for power generation and for heating. In some countries with more limited domestic pipeline infrastructure, LNG is loaded into trucks at the import terminal and then transported via truck to the location where the natural gas will be used, though this is still fairly uncommon and accounts for less than 5% of total LNG imported globally.³

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

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

A key institutional feature of the LNG industry is that the majority of trade is carried out under long-term contracts signed between LNG suppliers and downstream buyers. A typical long-term contract specifies the average annual contracted quantity to be sold by the seller to the buyer, the start and the end date, and a pricing formula

³LNG trucking has been most extensively used in China, which accounted for 74% of trucked LNG globally in 2017.

that is used to determine the price under which trade takes place. Typically the price of contracted LNG is indexed to the price of some benchmark. The price of crude oil is the most common such benchmark, while for terminals operating out of the US, the price is typically indexed to the domestic price of natural gas in the Henry Hub market.

In addition to these basic features, a contract usually includes several additional clauses. The contract may specify a “take-or-pay” share: if the take-or-pay share is $X\%$ and the quantity agreed is Q , then the buyer commits to taking at least $X\%$ of Q every year; if the buyer fails to take delivery of the $X\%$ of Q in a given year, they must nevertheless pay for that LNG. This ensures that the seller’s minimum revenue from the contract is the price multiplied with the take-or-pay quantity. A contract may include a destination clause, which prohibit buyers from reselling the product outside a pre-defined market (typically the buyer’s home country) (IGU, 2009). Finally, the contract may also include a diversion clause which specifies how the two parties split the surplus in the event that they decide to sell the cargo to a third party (which is known in industry parlance as a “diversion”).

Both destination clauses, and to a lesser extent diversion clauses, have been controversial in the LNG industry. Destination clauses directly limit the extent to which buyers can engage in arbitrage. As a result, destination clauses have been challenged by anti-trust authorities in LNG importing countries. The European Union Commission ruled destination clauses anti-competitive in 2003 (IEA, 2013), following which they were gradually phased out of LNG contracts signed by European buyers.⁴ In 2017, the Japanese Fair Trade Commission prohibited destination clauses in new LNG contracts (Harding and Sheppard, 2017). Diversion clauses, by contrast, have not been directly prohibited, but they can also serve to reduce the incentives of buyers to engage in

⁴As Talus (2011) documents, negotiations between the EU commission and LNG suppliers over the removal of destination clauses lasted for several years.

arbitrage, since they reduce and in some cases eliminate the profits the buyer earns from arbitrage (which now have to be shared with the seller).⁵

Figure 2.1 shows the distribution of contract duration for long-term contracts (defined for the purposes of this study as contracts that are longer than 4 years in duration). The majority of contracts are well over 10 years in length, with the modal contract length being around 20 years.

Long-term contracts do not provide the only way for parties to trade LNG, however. Parties can also trade LNG on the spot market, or using short-term contracts. A typical spot transaction involves the delivery of a single LNG cargo from the seller's export terminal to the buyer's import terminal.⁶ Figure 2.2 shows that the share of spot and short-term trade has been rising over time, accounting for only 12% of trade in 2004 but 27% of trade in 2017. Finally, buyers can also "re-export" LNG to other buyers by purchasing LNG from one source and re-loading the LNG onto a new tanker, though this accounts for a relatively small proportion of overall trade.⁷

Historically, financial markets in LNG have been very limited in size and scope, in stark contrast to the global crude oil market or the domestic natural gas market in the United States. Since 2017, though, there has been increasing trade in LNG derivatives. Derivatives represented only about 2% of LNG trade volumes at the beginning of 2017,

⁵Talus (2011) distinguishes between two kinds of diversion clauses: "profit-splitting" clauses that split the eventual profits from resale between the buyer and seller, and "price-splitting" clauses that split the resale price. He argues that "price-splitting" clauses often entirely remove the buyer's incentive to resell, as they will sometimes have to do so at a loss.

⁶Each LNG tanker will carry "one cargo". The volume of LNG carried by a fully loaded LNG tanker is on average around 132,000 m^3 , though it differs depending on the size of the tanker. There were a total of 5019 cargoes delivered in the year 2017.

⁷The share of re-exports increased from 0.15% in 2008 to a peak of 2.69% in 2014, but decreased to 1% of total trade by 2017.

but by the end of 2018, the share had grown to around 23% (Stapczynski and Murtagh, 2019). It should be noted that the size of the derivatives market is still small relative to the physical market: for comparison, in the crude oil market, derivatives volumes account for around 17 times the volume of physical trade (Terazono, 2019). This paper focuses on the LNG market up until 2017, when derivatives trade played only a very limited role in the market.

Figure 2.3 shows the evolution of industry capacity and trade over time. It plots total liquefaction capacity (which is the capacity of LNG exporters), total regasification capacity (which is the capacity of LNG importers) and total LNG trade. The binding constraint on the volume of trade tends to be the available liquefaction capacity, and export capacity utilization tends to be high, ranging from between 80 to 90%. By contrast, there is generally a lot of excess regasification capacity, and this has grown in recent years with increasing entry of new importing countries. These patterns reflect the fact that liquefaction projects are typically significantly more costly to build than regasification projects: the capital cost of a liquefaction project is on average equal to \$1.7 billion per million tonnes per annum (MTPA) of capacity, whereas the capital cost of a regasification project is on average \$250 million per MTPA (OIES, 2017).


Data

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

Data on LNG contracts I observe data on individual LNG contracts. This data is available from Bloomberg.⁸ I corroborate this with data with annual industry reports


⁸The dataset was accessed through the Bloomberg terminal.

Figure 2.1: Histogram of long-term contract durations



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


Figure 2.2: Share of short-term contracts and spot transactions in LNG trade



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

Source: GIIGNL Annual Reports.

Figure 2.3: Liquefaction capacity, regasification capacity and LNG trade over time



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

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

For each contract, I observe the contract quantity, year during which contract was

signed, and the contract duration (i.e. the contract start year and end year).⁹ I observe the identity of the buyer, the identity of the seller, the import country and the export country. About 10% of the contracts (mostly in the latter part of the sample) are “portfolio” contracts where the seller is free to supply the LNG from anywhere in the world where it can gain access to LNG; in such cases the export country is recorded as “Unspecified”. I observe whether the contract is signed on a free-on-board (FOB) basis (which loosely speaking means that the buyer is responsible for the shipping) or a delivered ex-ship (DES) basis (which means that the seller is responsible for the shipping). The dataset also records whether the contract is an extension of an earlier contract, or whether the contract is part of a “swap” arrangement with another contract.¹⁰

There are a number of contract details that I do not observe. The most important of these is the pricing formula, which indexes the price of LNG to the price of a selected energy benchmark (such as the oil price or the domestic price of natural gas). The pricing formula is typically confidential and known only to the parties that are signatory to the contract. I also do not observe the “take-or-pay” share of the contract, whether the contract includes a destination restriction (which as described above restricts the ability of the buyer to resell LNG), and any “diversion” clauses that specify how the gains from “diverting” LNG cargoes to a third party are to be divided between the buyer and the seller.

Data on liquefaction and regasification capacity I obtain data on investment and capacity data, for both the liquefaction terminals (owned by exporters) and regasification

⁹Note that contracts are typically signed an average of three years prior to the period when trade begins. For an LNG contract beginning in 2010, for example, the buyer and seller would typically sign a sale and purchase agreement by 2007 or 2008.

¹⁰An example of a swap is when buyer A originally has a contract with seller A', and buyer B originally has a contract with seller B'. A swap deal would mean that buyer A now purchases from seller B', while buyer B now purchases from seller A'. This kind of arrangement is relatively rare in the industry.

terminals (owned by importers), from the annual reports of the GIIGNL (The International Group of Liquefied Natural Gas Importers), whose members collectively own nearly every LNG import terminal in the world. The dataset covers all liquefaction and regasification terminals operational from 2004 to 2017. The dataset consists of the start-up year, nameplate capacity, the ownership structure and the operator for every export and import terminal built on or prior to 2017. I also observe the year when a Final Investment Decision (FID) is made on an export terminal¹¹. Due to time-to-build, the FID year is earlier than the start-up year (which is when the terminal begins operating).

Data on LNG trade flows I obtain data on LNG trade flows from the annual reports of the GIIGNL. This dataset reports yearly LNG trade flows for each country pair from 2004 to 2017. We separately observe trade flows that take place under long-term contracts (exceeding four years in length), short-term contracts and spot trade¹², and re-exports.¹³ A novel feature of this dataset is that it allows me to distinguish between flows that take place under long-term contracts and short-term/spot contracts. Actual trade flows under long-term contracts may differ from the quantity originally agreed to in the contracts between the two countries, either due to production fluctuations or because the parties to the contract agreed to divert cargoes to other destinations.¹⁴

LNG spot prices and freight costs I obtain data on weekly LNG spot prices and

¹¹Typically the Final Investment Decision (FID) is preceded by various feasibility studies, obtaining the relevant regulatory permits and signing sale and purchase agreements or contracts with buyers. The FID itself is the decision by all project partners to finally commit to the project. Construction of an LNG export terminal only begins once a Final Investment Decision has been taken by the investors.

¹²Short-term contracts are defined as contracts that are four year or shorter in duration. The dataset does not separately distinguish short-term contracts from one-time spot transactions.

¹³This records exports from one importing country to another importing country, which is distinct from exports from an exporting to an importing country.

¹⁴It should also be noted that this dataset record country-country trade flows, rather than firm-firm trade flows.

shipping costs between February 2006 and August 2018 from several sources. The most comprehensive of these datasets comes from Waterborne Energy¹⁵. 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.¹⁶

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

Table 3.1 contains summary statistics on key variables used in the analysis. Trade flows and shipping costs are defined at the exporter-importer-year level; export capacity and total exports are defined at the exporter-year level and spot prices and total imports are defined at the importer-year level. A fact to note is the prevalence of zeros in the trade data: out of 6,406 observations at the exporter-importer-year level, only 1,563 observations feature positive spot trade flows and 1,208 observations feature positive long-run contracted flows.

¹⁵I accessed this dataset through the Reuters Eikon terminal.

¹⁶Note that the dataset does not include freight rates for every possible importer-exporter pair. Freight rates for exporter-importer pairs not covered by Waterborne are imputed based on a regression model linking the freight rate to the distance between two ports.

Table 2.1: Summary Statistics

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

¹. All trade variables (spot trade, contracted trade, total exports, total imports etc.) are measured in million tonnes or mt. Capacity is measured in million tonnes per annum, or mtpa.


². Spot prices and shipping costs are measured in US\$/MMBtu.

2.3 Descriptive Evidence

I now discuss some of the features of the industry that are particularly relevant for the analysis undertaken in this paper. Firstly, Figure 2.4 plots a histogram of the gap between the contract signature date and the final investment decision (FID) date. As we can see, a large share of contracts are signed prior to the final investment decision date. Moreover there appears to be a clustering of contract decisions in the months prior to and during the same month as the investment decision. Relatively fewer contracts are signed in the months right after the investment decision has been made. This data pattern is consistent with a hold-up story where sellers prefer to sign contracts with buyers before committing to the investment. Figure 2.5 shows that on average, a seller signs contracts amount to slightly over 60% of their capacity before they commit to the investment.


Sellers that are located far away from potential buyers are in a relatively weak

Figure 2.4: Gap between contracting year-month and investment year-month



Note: This chart shows the number of months between the contract signature date and the date at which sellers make a final investment decision. Negative values indicate that the contract was signed prior to investment; positive values indicate the contract was signed after the investment decision was made.

Figure 2.5: Cumulative share of capacity signed under long-term contracts



Note: This chart plots the average share of capacity signed under long-term contracts.

bargaining position and face greater hold-up risk, since they would find it difficult to find a replacement for their buyer. According to the theory developed later, we would expect such sellers to engage in a greater degree of pre-investment contracting compared to sellers that are located close to potential buyers. I test this prediction of the theory by regressing the share of contract quantities signed before the seller has committed to the investment on the average distance from the seller to LNG buyers. As Table 2.2 shows, the coefficient on average distance is positive and significant and is robust to controlling for the capacity of the liquefaction project, the share of capacity that is contracted and regional dummies. The size of the coefficient suggests that increasing the average distance between an export project and buyers by 1000 nautical miles will increase the share of contract quantity signed pre-investment by 11-16 percentage points.


Next I present evidence suggesting both that there is misallocation in the LNG industry, and that contract rigidities contribute to this. First, I describe the recent evolution of spot prices and spot trade flows in the industry, and argue that these patterns are difficult to reconcile with competitive allocation of LNG. Figure 2.6 shows spot prices in various LNG importing regions. The feature that stands out from this figure is the presence of significant and systematic spot price differentials across regions, especially during periods when the LNG market is tight. A particularly striking illustration of this is the period between mid-2011 and end-2013, when there was a large spike in Japan's LNG demand following the Fukushima nuclear disaster. Asian spot prices, as well as spot prices in Latin America, remained an average of \$5/MMBtu higher than European prices during this period. The prices converge again only in late 2014 and early 2015, driven by a combination of reduced LNG demand in Asia, lower transportation costs (driven by the oil price collapse) and addition of new export capacity starting from 2015. In addition to the post-Fukushima period from 2011 to 2013, there have been other periods when prices in different regions sharply diverged from one

Table 2.2: Regression of share of contract volume signed before final investment decision date on characteristics of the export project

	(1)	(2)	(3)	(4)
	Dependent variable: Share of contract quantity signed before final investment decision			
Average distance from buyers	0.11** (0.053)	0.12** (0.052)	0.14*** (0.048)	0.16** (0.077)
Capacity		-0.023* (0.013)	-0.026** (0.012)	-0.025* (0.013)
Share of capacity contracted			0.48*** (0.14)	0.43*** (0.13)
Atlantic				-0.43 (0.52)
Middle East				-0.59 (0.41)
Pacific				-0.69 (0.55)
N	58	58	58	58
R ²	0.068	0.12	0.28	0.36
Mean of dependent variable:	0.68			
Mean of average distance from buyers:	6.80			

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

Figure 2.6: LNG Spot Prices in Different Regions



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

Figure 2.7: LNG imports under long-term contracts to different regions





Figure 2.8: Short-run and spot LNG imports to different regions



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

another. For example, between January 2007 and July 2008, another period of tight demand, Asian spot prices were about \$3.5/MMBtu higher on average than European spot prices.

In a competitive spot market with no frictions from either contracts or market power, faced with these divergent prices, capacity-constrained sellers would sell only to the destination with the highest price net of transportation costs. This would suggest that during these periods of high Asia-Europe spot price differentials, spot exports should be mostly directed to Asia, as even after accounting for transportation costs most sellers received higher prices from selling to Asia than to Europe. As Figure 2.8 shows, though, during the “boom” period of 2011-14 Europe continued to import a significant amount of LNG on the spot market. Moreover, much of Europe’s spot market imports during this period came from sellers that would have received higher prices from selling to Asia. For example, in 2012, 46% of Europe’s spot market purchases were from Qatar, and a further 12% were from Egypt and Algeria. This is in spite of the fact that these countries faced roughly similar shipping costs of selling LNG to Asia and Europe, meaning the marginal cargo garnered a premium of \$2-4/MMBtu if it were sent to Asia rather than Europe. The fact that sellers in Middle East and North Africa continued to sell to Europe is inconsistent with competitive behavior on the part of LNG exporters.

During the same period, Europe continued to import large amounts of contracted LNG, suggesting that long-term contracts also impede the market’s response to demand shocks. Figure 2.7 breaks down LNG imports under long-term contracts by different importing regions. After the Fukushima crisis, European countries continue to import significant amounts of contracted LNG, despite having a lower willingness-to-pay (on the margin) than Asian buyers (as indicated by the spot price differentials). Most of these contracts were signed well before the nuclear disaster in Japan. If exporters had

fewer contractual commitments to European buyers during this period, they would likely have sold more LNG to Asia and less LNG to Europe, resulting in a more efficient allocation of LNG. Further evidence that contracts impeded efficiency during this period is provided by re-sale data: during the 2011 to 2014 period, European re-sales to Asia rose eight-fold relative to the 2009-2010 period, before dropping sharply from 2014 to 2015 (when the price differentials between the two regions closed due to increased capacity and dampened demand).

The pattern of sellers selling to multiple markets at different received spot prices suggests that sellers deviate from perfectly competitive behavior. I now present a reduced-form test for whether contract-induced rigidities can cause sellers to deviate from efficient behavior. In an efficient market, sellers should sell most of their output to the destination from which they receive the highest price (net of shipping costs). However this might not happen if a large proportion of a seller's capacity is tied to long-term contracts that cannot be adjusted flexibly in response to demand changes. To test if contract rigidities can lead to misallocation, I test whether sellers with a large share of committed capacity are less responsive to price than sellers with mostly uncommitted capacity. Formally, I look at how the relationship between the share of a seller's LNG output that it allocates to destination j and the price that the seller receives from destination j depends on the seller's total share of committed capacity ($share_{it}^c$, which equals 0 if the seller only trades on the spot market and equals 1 if the seller sells all of its capacity under contracts).

$$\ln(q_{ijt}/q_{it}) = \gamma_0 + \gamma_1 \ln(p_{ijt}) + \gamma_2 \ln(p_{ijt}) * \ln(share_{it}^c) + X_{ijt} * \gamma_3 + \mu_{ijt}$$

The coefficient γ_1 captures the baseline responsiveness of sales to price, while γ_2 measures how the responsiveness changes depending on the share of the seller's capacity committed to long-term contracts. If long-term contracts were as flexible in responding

to demand as spot trade, then γ_2 ought to be 0: sellers with mostly committed capacity should be just as responsive to demand changes as sellers with mostly uncommitted capacity. Table 2.3 however shows that γ_2 is significantly negative, meaning that sellers with a lot of contracted capacity are less responsive to spot prices (and thus allocate their output less efficiently).

Table 2.3: Reduced Form Evidence for Contracting Frictions: Regression of Share of LNG Sales to market j on Price at market j

	(1)	(2)	(3)	(4)	(5)
	Dependent variable: $\log(\text{share of LNG sales to market } j)$				
$\gamma_1 \log(p_{ijt})$	0.35*** (0.078)	0.32*** (0.081)	0.55*** (0.11)	0.31** (0.12)	0.24*** (0.091)
$\gamma_2 \log(p_{ijt}) \log(\text{share}_{it}^c)$		-0.051 (0.034)	-0.094*** (0.035)	-0.17*** (0.054)	-0.20*** (0.031)
Year FE			Y	Y	Y
Exporter FE				Y	Y
Exporter-Importer FE					Y
N	1737	1737	1737	1737	1737
R^2	0.012	0.013	0.041	0.14	0.82

Note: Each observation is an exporter-importer-year pair. Total LNG exports from country i to country j in year t are measured in million tonnes. The price i receives from j , p_{ijt} , is measured in \$/MMBtu. The variable share_{it}^c refers to the proportion of seller i 's total capacity that is sold under long-term contracts. Statistical significance at the 10%, 5%, and 1% levels are denoted with *, **, and ***, respectively.

2.4 Model

2.4.1 Overview of model

In this section, I describe a model of contracting, investment and spot trade in the LNG industry. The model features a set of sellers and buyers that are located in geographically distinct locations. Buyers have uncertain demand for LNG. Sellers invest in LNG export projects that allow them to convert natural gas to LNG, which they ship to buyers. Trade takes place via both long-term contracts as well as a spot market.

The model consists of two components. The first component is a model of contracting and investment decisions. Sellers decide how much capacity to build. A seller and a buyer can sign a long-term contract which specifies a fixed quantity to be traded between the buyer and the seller each year, over an agreed contract duration, as well as a lump-sum transfer to be paid by the buyer to the seller. The choice of the long-term contract quantity and lump-sum transfer is the outcome of a Nash bargaining game played between the seller and buyers. Importantly, long-term contracts can be signed either before or after the seller commits to the investment. The timing has economic consequences because when contracts are signed prior to investment, the cost of the investment is not yet sunk, as opposed to contracts that are signed post-investment.

The second component is a model of the LNG spot market. The spot market takes place every year. On the spot market, both capacity and contracted quantities are taken as given. Demand shocks are realized and sellers then decide how much LNG to sell to each buyer on the spot market. Spot market flows are modelled as a Cournot game played among the sellers.

The model features the key economic mechanisms underlying contractual choices.

Firstly, because the seller bargains with the buyer over how to split the surplus from contracting, the model allows for both the possibility of hold-up. If buyers and sellers could only bargain *after* investments are made, and if buyers have some bargaining power, then sellers will under-invest (since they do not recoup the full marginal value of their investment). Likewise, sellers who invest without signing any contracts and only sell on the spot market also face hold-up, since the cost of investment is already sunk by the time spot trade takes place. However, the ability to sign contracts prior to investing allows sellers to partly overcome the hold-up problem, since when bargaining prior to investing, the seller has the credible threat of not investing as a way to push up the price that it receives from the buyer.

Secondly, the model incorporates the effects of long-run contracts on the short-run allocation of LNG. Because contract quantities are fixed and have to be agreed before demand shocks are realized, the use of long-term contracts can reduce the flexibility of the market in responding to demand shocks, which leads to welfare losses from short-run misallocation. At the same time, because the model allows for sellers to exercise market power on the spot market, under some conditions contracts can also lead to improved short-run allocation of LNG as highlighted by Allaz and Vila (1993) and related papers.

I describe the model of the spot market in Section 2.4.2, which features the “static” decisions (i.e. spot market trade flows and demand). This section draws heavily from the first chapter of my dissertation, but takes as given the fact that sellers exercise market power. Following that, I describe the model of long-run contracting and investment in Section 2.4.3.

2.4.2 Demand, production and spot trade

I begin by describing a model of LNG spot market flows, which takes place every year. By this stage of the game, all investment decisions have been made and long-run contracts already committed to. Buyer demand shocks are now realized. Sellers then choose how much to sell on the spot market.

Setup

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 indexed by $i = 1, \dots, N$. Buyers and sellers are risk-neutral and have discount factor β .

In the empirical analysis, each exporting country is treated as a separate seller, and each importing country is treated as a separate buyer.

Buyers and sellers are geographically differentiated from one another, and d_{ij} is the distance between seller i and buyer j . Buyer and seller attributes are denoted by x_{jt} and x_{it} respectively.

Demand

Buyers have demand curves $Q_d(p_{jt}, x_{jt}, \varepsilon_{jt})$, where p_{jt} is the spot price. The demand curve for buyer j is given by the following equation:

$$Q_{jt} = Q_d(p_{jt}, x_{jt}, \varepsilon_{jt}) \quad (2.1)$$

where ε_{jt} denotes demand shocks, p_{jt} denotes the spot price paid by buyer j and x_{jt} denotes demand shifters. The inverse demand curve is $p_{jt} = P_j(Q_{jt}, x_{jt}, \varepsilon_{jt})$. Note that

Q_{jt} denotes buyer j 's *total* purchases of LNG, including both purchases under long-term contracts and purchases on the spot market.

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

$$Q_{jt} = -bp_{jt} + \theta_j + x_{jt}\theta_{dx} + \varepsilon_{jt} \quad (2.2)$$

Here θ_j denotes country fixed effects.

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. The total amount of LNG produced by seller i equals $q_{it} = \sum_j q_{ijt} = \sum_j (S_{ijt} + q_{ijt}^c)$, while the total quantity imported by buyer j equals $Q_{jt} = \sum_i q_{ijt} = \sum_i (S_{ijt} + q_{ijt}^c)$.

Contracted quantities q_{ijt}^c are pinned down by long-term contracts and cannot adjust in the short-run in response to the demand shocks. This reflects the empirical reality in the LNG industry that most contracts are signed well in advance of when trade takes place, and have features such as destination clauses that limit the flexibility of the quantity traded. In contrast, S_{ijt} can be chosen by the sellers every year after observing the demand shocks, and is the key strategic variable in the model of the spot market.

Costs of production and sales

Sellers incur costs both in producing LNG, and then in selling the LNG to the buyers. Let $C(q_{it}, K_{it})$ denote the cost of production, where K_{it} is seller i 's total capacity in period t . The inclusion of K_{it} in the cost function reflects the fact that capacity constraints are

important in this industry (as evidenced by Figure 2.3) and sellers with larger capacity are able to export more LNG.

In the empirical application, I assume that firms face a hockey-stick cost function where the marginal cost becomes steep as the firm nears full capacity utilization. This follows Ryan (2012) and Miller and Osborne (2014), who use a similar specification of the cost function in their analyses of the US portland cement industry. The firm faces constant marginal costs until its capacity utilization hits a threshold of v . If the firm's capacity utilization exceeds v , the firm's marginal costs increase with the level of capacity utilization in excess of the threshold, with the parameter δ_2 governing the rate at which marginal costs increase as the firm approaches full capacity utilization:

$$C(q_{it}, K_{it}) = \delta_1 q_{it} + \frac{\delta_2}{2} 1(q_{it} \geq vK_{it})(q_{it} - vK_{it})^2 \quad (2.3)$$

In addition, the firm incurs costs of trading that differ by buyer, which consists of two components. Each unit of LNG costs c_{ijt}^d to ship from i to j , where the shipping cost c_{ijt}^d is increasing in the distance d_{ij} between seller i and buyer j and fluctuates by year.¹⁷ In addition, the costs of selling LNG to buyer j can also depend on the characteristics of the buyer and seller or x_{ijt} . For example, there may be costs of building a relationship which are only incurred when a seller and buyer first engage in trade; in this case, x_{ijt} would include an indicator for whether or not i and j have traded in the past.

Thus firm i 's total costs from LNG sales of $\{q_{ijt}\}_{j=1}^J$ are given by:

$$C_i(\{q_{ijt}\}_{j=1}^J, K_{it}) = \underbrace{C(q_{it}, K_{it})}_{\text{Cost of production}} + \underbrace{\sum_j (c_{ijt}^d + \phi(x_{ijt}))q_{ijt}}_{\text{Cost of trading}} \quad (2.4)$$

¹⁷ c_{ijt}^d is directly observed in the dataset. In practice c_{ijt}^d is a linear function of the nautical distance d_{ij} between seller i and buyer j , since the cost of shipping increases linearly with the journey time. The slope of this linear function as well as the intercept term differ by year, reflecting prevailing LNG shipping rates at the time.

Spot market equilibrium

Each of the J markets clears separately and simultaneously. The market clearing price vector $p_t^* = (p_{1t}, \dots, p_{jt}, \dots, p_{Jt})$ is characterized by the following set of equations:

$$Q^d(p_{jt}^*, \epsilon_j) = \sum_{i=1}^N q_{ijt}^c + \sum_{i=1}^N S_{ijt}, \forall j \quad (2.5)$$

I model the spot market equilibrium as a Cournot game played between the N sellers. In a Cournot equilibrium, each seller i takes as given rival spot quantities $\{S_{-ijt}\}_{j=1}^J$ and chooses the vector of spot quantities, $\{S_{ijt}\}_{j=1}^J$, that maximizes its sum of profits across all markets:

$$\{S_{ijt}\}_{j=1}^J = \underset{\{\hat{S}_{ijt}\}_{j=1}^J}{\operatorname{argmax}} \left[\sum_{j=1}^J p_{jt}^*(S_{ijt}, S_{-ijt}) \hat{S}_{ijt} - C(q_{it}, K_{it}) - \sum_j (c_{ijt}^d + \phi(x_{ijt})) q_{ijt} \right] \quad (2.6)$$

where recall that $q_{it} = \sum_j (q_{ijt}^c + S_{ijt})$.

Notice that the seller's payoff function does not include revenue from contracted sales, since this is unaffected by spot market decisions and is effectively “sunk” by the time the seller reaches the spot market.

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

$$\underbrace{p_{jt}^* + S_{ijt} \frac{\partial p_{jt}^*(S_{ijt}, S_{-ijt})}{\partial S_{ijt}}}_{\text{Marginal revenue of selling to market } j} - \underbrace{\left(\frac{\partial C(q_{it}, K_{it})}{\partial S_{ijt}} + c_{ijt}^d + \phi(x_{ijt}) \right)}_{\text{Marginal cost of selling to market } j} \leq 0 \quad (2.7)$$

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

Payoffs to the buyers and sellers

The sellers take contracted quantities, capacity and demand shocks as given when choosing spot market quantities. Thus the seller i 's total profits in period t , as well as buyer j 's consumer surplus in period t , are functions of q_t^c (which is a vector of all contracted trade flows), K_t (a vector of capacities) and ε_t (a vector of demand shocks).

Integrating out the demand shocks, we can derive expected payoffs to sellers and buyers as functions of contracted quantities and capacities. Let $\pi_{it}^s(q_t^c, K_t)$ denote seller i 's expected payoff in period t . Similarly, let $\pi_{jt}^b(q_t^c, K_t)$ denote buyer j 's expected payoff (consumer surplus) in period t .

Discussion

I now discuss several assumptions of the spot market model and the resulting implications for the allocation and pricing of LNG on the spot market.

Because of capacity constraints, the seller's choice of S_{ijt} depends on how much they sell in other markets. If the seller sells high quantities in some other market (due to say a demand shock in that country), then $\frac{\partial C(q_{it}, K_{it})}{\partial S_{ijt}}$ will be higher because of increasing marginal costs and so the seller will tend to sell less in market j . This in turn implies that demand shocks in one market will affect firms' spot market sales across all markets. The *interconnectedness* of markets resulting from capacity constraints is an important feature of the model: if instead sellers faced constant returns to scale, their decisions in one market would be independent of their decisions in other markets.

The assumption that sellers engage in strategic competition is crucial since it implies that sellers engage in spatial price discrimination on the spot market. To see this consider

a seller i that sells to two buyers j and j' . For simplicity assume that the per-unit trade cost is identical whether or not the seller sells to buyer j or to buyer j' (e.g. this would happen if the seller is equidistant from both buyers and there are no non-shipping costs of trade). In that case, the first-order condition (2.7) implies that:

$$\underbrace{p_{jt}^* + S_{ijt} \frac{\partial p_{jt}^*(S_{ijt}, S_{-ijt})}{\partial S_{ijt}}}_{\text{Marginal revenue of selling to buyer } j} = \underbrace{p_{j't}^* + S_{ij't} \frac{\partial p_{j't}^*(S_{ij't}, S_{-ij't})}{\partial S_{ij't}}}_{\text{Marginal revenue of selling to buyer } j'}$$

Because the seller equalizes marginal revenue (and not price) across buyers, in general the seller will charge different prices to different buyers. As described in the first chapter of this dissertation, I have tested the Cournot model of firm behavior against the alternative hypothesis of perfect competition, by estimating a conduct parameter that captures the extent to which sellers deviate from perfectly competitive behavior. I find that seller behavior is consistent with the Cournot model and unlikely to be generated by the competitive model.

I assume that buyers cannot engage in arbitrage on the spot market. If costless arbitrage were possible, price differentials across regions would be arbitrated away, and sellers would not be able to engage in third-degree price discrimination. There are a number of barriers to arbitrage in LNG. The most important of these are resale restrictions that are commonly written into long-term LNG contracts. As discussed earlier, these take the form of destination clauses, which prohibit resales by buyers, and profit-sharing clauses, which require that profits from any LNG diversions to an alternative market be split between the seller and the buyer, reducing the incentive of the buyer to engage in arbitrage. Shipping costs constitute another major barrier to arbitrage. A buyer would find it profitable to re-export LNG to another buyer only if the price difference exceeds the shipping cost. As such in practice, re-exports in LNG are limited in volume and accounted for only about 1% of total trade volumes between

2004 and 2017¹⁸. Finally, financial markets for LNG have been historically very limited, and derivatives trade is small in relation to physical trade, further limiting the scope for arbitrage.

Throughout I also assume that buyers are price-takers when they purchase LNG on the spot market. Oligopsony is unlikely to be a first-order concern on the spot market, because there are many more importing countries than exporting countries, and the buyers' side of the market tends to be much more fragmented: there are usually several importing firms in each importing country.

2.4.3 Contracting and Investment

We now embed this model of short-run LNG flows into an equilibrium model of long-run contracting and investment.

Setup and Notation

Each seller i owns an LNG 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 exporting LNG.¹⁹

Long-term contracts can be signed either before or after the seller decides how much capacity to build. Let $\mathbf{B}_i^1$ denote the set of buyers with whom seller i can sign contracts before committing to investment (in Stage 1 of the game). Let $\mathbf{B}_i^3$ denote the set of

¹⁸Calculations based on data from GIIGNL's Annual Reports, 2004 - 2017.

¹⁹For notational simplicity, when presenting the model, I will abstract away from the possibility that a single seller may make multiple, separate investments. When estimating the model, I account for this possibility.

buyers with whom the seller can sign contracts after committing to investment (in Stage 3 of the game). These sets of buyers may partially overlap.

A long-term contract signed between seller i and buyer j consists of a start date t_{ij}^{begin} , an end date t_{ij}^{end} , a contracted quantity q_{ij}^c (to be delivered by the seller to the buyer for every year during which the contract is operational) and a lump-sum transfer T_{ij} which the buyer pays to the seller. Thus in a given year t , the total contracted quantity seller i delivers to buyer j , q_{ijt}^c , is the sum of contracted quantities across all contracts signed between seller i and buyer j that are still active in period t .

Let $q_i^{c,1} = \{q_{ij}^c\}_{j \in \mathbf{B}_i^1}$ denote a vector comprising all the contract quantities signed by seller i in Stage 1.

Similarly, let $q_i^{c,3} = \{q_{ij}^c\}_{j \in \mathbf{B}_i^3}$ denote a vector comprising the contract quantities signed by seller i in Stage 3.


Timing of the game

I model contracting and investment decisions for each seller i (and associated buyers) as a sequential, three-stage game, as summarized in Figure 2.9.

Stage 1 of the game takes place before any investments have taken place. Each seller-buyer pair then bargains over the contract quantity and a lump-sum transfer to be paid by the buyer to the seller.

In Stage 2, the seller chooses how much to invest, taking as given any contracts that have already been signed in Stage 1. The seller must build enough capacity to meet these pre-committed quantities.

Figure 2.9: Stages of the game



Stage 3 of the game occurs after the seller has committed to the investment. Just as in Stage 1, each seller-buyer pair then bargains over the contract quantity and a lump-sum transfer to be paid by the buyer to the seller.

Once the three stages of decision-making are complete, the seller completes construction of the project. Once construction is complete and the project becomes operational, the capacity from that project becomes available to the seller. The seller participates in the LNG market every year, meeting their contracted obligations and selling any uncommitted capacity on the annual spot market.

The bargaining model I use is based on Chipty and Snyder (1999). At each of the two contracting stages (Stages 1 and 3 in Figure 2.9), each seller-buyer pair first chooses a contract quantity to maximize their joint surplus from contracting, taking as given the contract quantity chosen by other seller-buyer pairs. They then choose 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. A key difference from Chipty and Snyder (1999) is that the bargaining affects the sellers' investment

decisions: contracts signed prior to investing directly affect investment (as the seller must build enough capacity to meet contracts) and contracts signed post investment affect sellers' incentives to invest in anticipation of the bargaining.

The only decisions I explicitly model are the decisions by the buyers and sellers on the quantity to be contracted and the lump-sum transfer, and the investment decisions made by the sellers. The exact timing of decisions (e.g. the particular year when a contract is signed or an investment committed to), the start and end date of the contracts, and the set of buyers the seller can negotiate with at each stage, are treated as exogenous. In other words, I do not explicitly model the timing of decisions or the matching between buyers and sellers. I discuss the implications of this later.

Beliefs and expected payoffs

The payoffs to a given seller or a buyer depend on the contracting and investment decisions taken by all sellers and buyers, as well as the realization of demand shocks. Thus it is important to specify what agents believe about rival actions and about the evolution of demand uncertainty in order to work out their actions at each stage of the game.

A key assumption I make is that shocks to demand, ε_{jt} , are identically and independently distributed (i.i.d.). This implies that agents do not learn about demand with time. As such, agents can perfectly foresee the future contracting and investment decisions of rival agents. While the assumption of no serial correlation in demand is strong, it aids in making the model tractable to estimate since we do not need to estimate agents' beliefs about rival future actions. I discuss later how to extend the analysis to handle the case when demand shocks are persistent.

Under the assumption of i.i.d. demand shocks, we can now work out the expected payoffs to a particular seller i and the buyers that seller i signs long-term contracts with. Let $Y_{-i} = (q_{-i}^c, K_{-i})$ be a vector capturing the contracting and investment decisions made in projects operated by sellers other than i . Seller i and the buyers with whom i contracts take Y_{-i} as given when making their own contracting and investment decisions.

Let t_i^3 denote the time when Stage 3 contracts are signed. At this point, all contracting and investment decisions pertaining to seller i 's project have been finalized. The expected payoffs to the buyer and seller are their discounted sum of lifetime expected payoffs from period t_i^3 onwards, as functions of their own contract quantities $\{q_i^{c,1}, q_i^{c,3}\}$ and capacity K_i , as well as contract capacities and capacities chosen by other sellers and buyers (which are captured in Y_{-i}). Seller i 's expected payoff, V_i^3 , can be written as:²⁰

$$V_i^3(q_i^{c,1}, q_i^{c,3}, K_i, Y_{-i}) = \sum_{t=t_i^3}^{\infty} \beta^{t-t_i^3} \pi_{it}^s(q_t^c, K_t) \quad (2.8)$$

where recall that $\pi_{it}^s(q_t^c, K_t)$ is seller i 's expected profits in period t .

Similarly buyer j 's expected payoff, v_j^3 , can be written as:

$$v_j^3(q_i^{c,1}, q_i^{c,3}, K_i, Y_{-i}) = \sum_{t=t_i^3}^{\infty} \beta^{t-t_i^3} \pi_{jt}^b(q_t^c, K_t) \quad (2.9)$$

where recall that $\pi_{jt}^b(q_t^c, K_t)$ is buyer j 's expected consumer surplus in period t .

Stage 3: contracting after investment

I now describe how decisions are made at each of the three stages of the game for each project, starting from the final stage or Stage 3. In Stage 3, the seller (who has already committed to building a capacity of K_i) can negotiate contracts with a set of buyers $\mathbf{B}_i^3$.

²⁰The superscript “3” is to indicate that these are their payoffs after Stage 3 is complete.

Contracts are negotiated via the Nash bargaining model: each seller-buyer pair picks the quantity that maximizes their joint surplus and divides the incremental surplus in proportion to their respective bargaining weights, taking as given that every other seller-buyer pair picks quantity and prices in the same way.

The expected lifetime payoffs to the sellers and buyers are given by equations (2.8) and (2.9). In addition, I allow for the possibility that buyers may be willing to pay a “premium” to purchase contracted LNG (over spot LNG). This captures any additional utility buyers get from purchasing LNG using long-term contracts as opposed to the spot market. There are several interpretations of this contract premium. By signing a long-term contract, buyers can avoid any transaction costs of repeatedly purchasing large amounts of LNG on the spot market. Buyers may also pay a premium for contracted LNG due to energy security concerns, where contracts allow the buyer to lock in a portion of their purchases. Additionally, if buyers are risk-averse, they may prefer contracts since that can reduce the volatility of the costs of purchasing LNG to the buyer.

Let $w_j(q_{ij}, \eta_{ij}^3)$ denote the contract premium the buyer receives from signing a contract of quantity q_{ij} with seller i , where η_{ij}^3 is a publicly observable shock to the marginal value of contracting between seller i and buyer j . We assume that $w_j(q_{ij}, \eta_{ij}^3)$ is separable in $\bar{w}_j(q_{ij})$ and $\eta_{ij}^3 q_{ij}$:

$$w_j(q_{ij}, \eta_{ij}^3) = \bar{w}_j(q_{ij}) + \eta_{ij}^3 q_{ij}$$

Equilibrium Quantities

Each 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 “Nash-in-Nash” quantities

are thus given by:

$$q_{ij}^{c,3} = \operatorname{argmax}_{q_{ij}} [V_i^3(q_i^{c,1}, q_i^{c,3}, K_i, Y_{-i}) + v_j^3(q_i^{c,1}, q_i^{c,3}, K_i, Y_{-i}) + w_j(q_{ij}, \eta_{ij}^3)] \quad (2.10)$$

$$\forall j \in \mathbf{B}_i^3$$

The FOC to the quantity problem is:

$$\frac{\partial V_i^3}{\partial q_{ij}} + \frac{\partial v_j^3}{\partial q_{ij}} + \frac{\partial \bar{w}_j(q_{ij})}{\partial q_{ij}} + \eta_{ij}^3 = 0 \quad (2.11)$$

Note that the contracted quantity signed in Stage 3 does not depend on the cost of investment, since investment is already sunk at this stage. In addition, the contracted quantity also does not depend on the bargaining weight, since in the Chitty-Snyder bargaining model, the seller-buyer pair maximize their joint surplus regardless of how bargaining power is distributed.

Equilibrium transfers

The buyer-seller transfers are determined via Nash-in-Nash bargaining. Each seller-buyer pair chooses a transfer to maximize the Nash product of their surplus from contracting, assuming that every other pair reaches agreement.

We assume that if negotiations fail between seller i and buyer j , they move on to the spot market, where the seller can sell his additional capacity and the buyer can source for his additional LNG. Thus the spot market constitutes the outside option for both the seller and buyer.

Let $V_i^3(q_i^{c,1}, q_{i,\setminus ij}^{c,3}, K_i, Y_{-i})$ denotes seller i 's disagreement payoff when negotiating with buyer j , where $q_{i,\setminus ij}^{c,3}$ denotes the vector of contract quantities if we set q_{ij}^c to 0 but maintain all other contracts in $q_i^{c,3}$. Similarly $v_j^3(q_i^{c,1}, q_{i,\setminus ij}^{c,3}, K_i, Y_{-i})$ denotes buyer j 's disagreement payoff when negotiating with seller i .

Then the lump-sum transfers T_{ij}^c that seller i receives from buyer j in Stage 3 are given by:

$$T_{ij}^c = \operatorname{argmax}_T \left(V_i^3(q_i^{c,1}, q_i^{c,3}, K_i, Y_{-i}) - V_i^3(q_i^{c,1}, q_{i,\setminus ij}^{c,3}, K_i, Y_{-i}) + T \right)^\tau \left(v_j^3(q_i^{c,1}, q_i^{c,3}, K_i, Y_{-i}) + w_j(q_{ij}^c, \eta_{ij}^3) - v_j^3(q_i^{c,1}, q_{i,\setminus ij}^{c,3}, K_i, Y_{-i}) - T \right)^{1-\tau} \quad \forall j \in \mathbf{B}_i^3 \quad (2.12)$$

where τ is the bargaining weight of seller i when negotiating with buyer j ²¹.

We can show that the Nash-in-Nash transfers equal:

$$T_{ij}^c = \tau \left(v_j^3(q_i^{c,1}, q_i^{c,3}, K_i, Y_{-i}) + w_j(q_{ij}^c, \eta_{ij}^3) - v_j^3(q_i^{c,1}, q_{i,\setminus ij}^{c,3}, K_i, Y_{-i}) \right) - (1 - \tau) \left(V_i^3(q_i^{c,1}, q_i^{c,3}, K_i, Y_{-i}) - V_i^3(q_i^{c,1}, q_{i,\setminus ij}^{c,3}, K_i, Y_{-i}) \right) \quad \forall j \in \mathbf{B}_i^3 \quad (2.13)$$

As one might expect, 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

We can now define the payoffs for sellers and buyers at the end of Stage 2 and just prior to the beginning of Stage 3.

The expected payoff for seller i at the end of Stage 2, V_i^2 , is his expected payoff at the end of Stage 3, V_i^3 , plus whatever transfers he receives from the buyers that he is matched to in Stage 3.²² The payoffs for any buyer j that contracts in Stage 3 is the expected payoff at the end of Stage 3, v_j^3 , plus the additional contract premium w , minus any transfers made to the sellers.

²¹It is straightforward to allow the bargaining weight to differ by seller and buyer.

²²The superscript "2" indicates these are payoffs after Stage 2 is complete.

$$V_i^2(q_i^{c,1}, K_i, Y_{-i}) = V_i^3(q_i^{c,1}, q_i^{c,3}, K_i, Y_{-i}) + \sum_{j \in \mathbf{B}_i^3} T_{ij}(q_i^{c,1}, q_i^{c,3}, K_i, Y_{-i}) \quad (2.14)$$

$$v_j^2(q_i^{c,1}, K_i, Y_{-i}) = v_j^3(q_i^{c,1}, q_i^{c,3}, K_i, Y_{-i}) + w_j(q_{ij}^c, \eta_{ij}^3) - T_{ij}(q_i^{c,1}, q_i^{c,3}, K_i, Y_{-i}) \quad (2.15)$$

Stage 2: investment

In Stage 2, the seller chooses how much capacity to build. The seller is already committed to any contracts that he has signed in Stage 1, and additionally takes into account contracts that he may sign in the future (in Stage 3) as well as future expected spot sales when deciding how much capacity to build.

As noted earlier, $V_i^2(q_i^{c,1}, K_i, Y_{-i})$ is the seller i 's expected payoff from constructing capacity equal to K_i . The superscript indicates again that this is the payoff the seller considers during Stage 2 of the game.

Let $C_i(K_i, \eta_i^2)$ denote the cost of the investment. η_i^2 is a publicly observable shock to the marginal cost of investing by seller i . We assume that investment costs are separable in the shock:

$$C_i(K_i, \eta_i^2) = \bar{C}_i(K_i) + \eta_i^2 K_i \quad (2.16)$$

The seller's optimal choice of K_i is given by:

$$K_i^* = \operatorname{argmax}_{K_i} [V_i^2(q_i^{c,1}, K_i, Y_{-i}) - C_i(K_i, \eta_i^2)] \quad (2.17)$$

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

$$\frac{\partial V_i^2}{\partial K_i} - \frac{\partial \bar{C}_i(K_i)}{\partial K_i} - \eta_i^2 = 0 \quad (2.18)$$

Payoffs at the end of Stage 1

We can now define the payoffs for sellers and buyers just prior to the beginning of Stage 2 and at the end of Stage 1. Let $V_i(q_i^{c,1}, Y_{-i})$ denote seller i 's payoffs for any set of contract quantities in Stage 1.²³ Let $v_j(q_i^{c,1}, Y_{-i})$ denote buyer j 's payoffs from a vector of contract quantities in Stage 1. These payoffs equal:

$$V_i^1(q_i^{c,1}, Y_{-i}) = V_i^2(q_i^{c,1}, K_i, Y_{-i}) - C_i(K_i, \eta_i^2) \quad (2.19)$$

$$v_j^1(q_i^{c,1}, Y_{-i}) = v_j^2(q_i^{c,1}, K_i, Y_{-i}) \quad (2.20)$$

Stage 1: contracting prior to investment

In Stage 1, just as with Stage 3, I allow for the buyer to receive a premium from signing a contract of quantity q_{ij} with seller i . The micro-foundation behind this is the same as that for Stage 3 payoffs: we can think of these as representing buyer preferences for energy security, risk aversion or avoiding transaction costs from repeated purchases on the spot market.

The contract premium term equals: $\omega_j(q_{ij}, \eta_{ij}^1)$, where each pair receives a publicly observable shock η_{ij}^1 that affects the marginal value from contracting. Just as with Stage 3, we assume that $\omega_j(q_{ij}, \eta_{ij}^1)$ is separable in $\bar{\omega}_j(q_{ij})$ and $\eta_{ij}^1 q_{ij}$:

$$\omega_j(q_{ij}, \eta_{ij}^1) = \bar{\omega}_j(q_{ij}) + \eta_{ij}^1 q_{ij}$$

Equilibrium Quantities

²³The superscript “1” indicates these are payoffs after Stage 1 is complete.

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

$$q_{ij}^c = \operatorname{argmax}_{q_{ij}} [V_i^1(q_i^{c,1}, Y_{-i}) + v_j^1(q_i^{c,1}, Y_{-i}) + \bar{\omega}_j(q_{ij}) + \eta_{ij}^1 q_{ij}] \quad (2.21)$$

$$\forall j \in \mathbf{B}_i^1$$

The FOC to the quantity problem is:

$$\frac{\partial V_i^1}{\partial q_{ij}} + \frac{\partial v_j^1}{\partial q_{ij}} + \frac{\partial \bar{\omega}_j(q_{ij})}{\partial q_{ij}} + \eta_{ij}^1 = 0 \quad (2.22)$$

Equilibrium Transfers

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

The lump-sum transfers T_{ij}^c that seller i receives from buyer j in Stage 1 are determined by Nash-in-Nash bargaining, with the equilibrium transfers characterized by the following equation:

$$T_{ij}^c = \operatorname{argmax}_T (V_i^1(q_i^{c,1}, Y_{-i}) - V_i^1(q_{i,\setminus ij}^{c,1}, Y_{-i}) + T)^\tau$$

$$(v_j^1(q_i^{c,1}, Y_{-i}) + \omega_j(q_{ij}^c, \eta_{ij}^1) - v_j^1(q_{i,\setminus ij}^{c,1}, Y_{-i}) - T)^{1-\tau}$$

$$\forall j \in \mathbf{B}_i^1 \quad (2.23)$$

where τ is the bargaining weight of seller i when negotiating with buyer j

2.4.4 Equilibrium properties of the contracting and investment game

I now explore some of the key properties of the model. I also discuss the implications of relaxing some of the assumptions that I make.

Holdup, bargaining and contracting

A first set of properties of the model relates to the relationship between holdup risk and contracting. Firstly, we show that sellers face holdup risk when investing. The lump-sum transfer that the seller receives in Stage 3 (post-investment contracting) depends on the seller's and buyer's respective gains from trade, as we can see from equation (2.14). Thus the seller's incentives to invest depend on the bargaining power of the seller relative to the buyer.

To see how this can generate a hold-up effect, 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 bilaterally optimal investment level) is:

$$K_3^{**} = \operatorname{argmax}_{K_i} [V_i^3(q_i^{c,1}, q_i^{c,3}, K_i, Y_{-i}) + v_j^3(q_i^{c,1}, q_i^{c,3}, K_i, Y_{-i}) + w_j(q_{ij}^c, \eta_{ij}^3) - C_i(K_i, \eta_i^2)] \quad (2.24)$$

By contrast, the actual investment chosen by the seller satisfies the following first-order condition (which can be derived by combining equations (2.14) and (2.16)):

$$K_i^* = \operatorname{argmax}_{K_i} [V_i^3(q_i^{c,1}, q_i^{c,3}, K_i, Y_{-i}) + \tau(v_j^3(q_i^{c,1}, q_i^{c,3}, K_i, Y_{-i}) + w_j(q_{ij}^c, \eta_{ij}^3) - v_j^3(q_i^{c,1}, q_{i \setminus ij}^{c,3}, K_i, Y_{-i})) - (1 - \tau)(V_i^3(q_i^{c,1}, q_i^{c,3}, K_i, Y_{-i}) - V_i^3(q_i^{c,1}, q_{i \setminus ij}^{c,3}, K_i, Y_{-i})) - C_i(K_i, \eta_i^2)] \quad (2.25)$$

Comparing equations (2.24) and (2.25), we can see that the seller's actual investment K_i^* will only reach the bilaterally optimal investment level K_i^{**} if two conditions are satisfied. First, the Nash bargaining weight τ has to equal 1, meaning that the seller has the ability to make take-it-or-leave-it offers. Secondly, the buyer's disagreement payoff, $v_j^3(q_i^{c,1}, q_{i,\setminus ij}^{c,3}, K_i, Y_{-i})$, has to equal 0. In other words, the seller invests at the optimal level only if they have the ability to fully extract the surplus from the investment. If instead the buyer has some bargaining power (i.e. if $\tau < 1$ or if the buyer has a non-trivial outside option), the seller will under-invest due to the holdup problem.


Secondly, the seller and buyer(s) have incentives to sign contracts in Stage 1, before the seller has committed to investing, as a way to reduce underinvestment from holdup. The cost of investment directly enters into V_i^1 , the expected payoff of the seller at the time of pre-investment contracting. This means that the transfers received by the seller in Stage 1 will account for the cost of investment: the higher the cost of investment, the higher the transfer received by the seller. In other words pre-investment contracting is not subject to the hold-up problem.

I carried out Monte Carlo simulations from a simplified version of the model to explore these properties. In the simplified version of the model, a single seller interacts with a single buyer, and both the seller and the buyer have access to a spot market (that forms their “outside option” in case the bargaining breaks down). The seller and buyer play the three-stage game outlined above: they can sign a contract prior to investment (Stage 1); the seller then chooses how much to invest (Stage 2); after the seller has already committed to the investment, the seller and buyer can get together again to sign a new contract, on top of the existing contract they already signed (Stage 3). Finally once contracts are signed and investments made, the seller sells any excess capacity on the spot market, and the buyer purchases any excess LNG requirement on the spot

market.

Consistent with the theoretical predictions, we find that if the seller and the buyer cannot contract in Stage 1, and as long as the buyer has some bargaining power, the seller will tend to under-invest relative to the optimum. This is the hold-up effect in action. The more bargaining power the seller has, the less severe the holdup and the higher the level of capacity the seller builds, as illustrated in Figure 2.10.


Figure 2.10: 1 seller, 1 buyer, no Stage 1 contracting: effect of seller bargaining power on K



Next, if we do allow the seller and buyer to contract before investment, we find that the ability to contract prior to investment increases investment and raises welfare. Moreover, the lower the Nash bargaining weight of the seller, the more contracting takes place in Stage 1 and the less contracting takes place in Stage 3 (see Figure 2.11). This is because the hold-up problem is more severe when sellers have low bargaining power, and anticipating this the seller and buyer sign larger volume contracts in Stage 1.

Finally the model also predicts that contracting decisions are affected by the availability of spot LNG. If the cost of purchasing spot LNG is high (e.g. because the buyer and seller are located faraway from other markets), then the buyer and seller

Figure 2.11: Share of contracts signed in Stage 1 (pre-investment) vs. seller bargaining power



will be more likely to engage in contracted trade. Moreover, if there is asymmetry in these outside options, the parties will adjust the extent to which they contract pre- and post-investment. For instance, if the buyer is located close to other sellers on the spot market and the seller is located far away from other buyers on the spot market, the seller's outside option is weaker than the buyer's, giving the seller a poorer bargaining position. This worsens the hold-up problem and the parties adapt to this by signing a higher contract quantity in Stage 1 and a lower contract quantity in Stage 3.

Long-term contracts and short-run allocative efficiency

How do long-term contracts affect short-run allocative efficiency? There are two competing forces at work in the model. On the one hand, long-term contracts are signed before demand is fully known, reducing the ability of firms to respond to ex-post demand shocks. This effect can lead to reduced allocative efficiency from the use

of contracts. 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 when firms have signed contracts, they 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 contracted output is unaffected by what happens on the spot market).

The allocative efficiency consequences of contracts hinges on which of these two competing effects dominates. Intuitively, we would expect the trade-off to depend on the extent to which firms are capacity-constrained. If firms produce well below capacity, contracts signed with one buyer do not hurt their ability to meet demand shocks experienced by other buyers; here, one would expect the Allaz and Vila effect to dominate and for contracts to improve short-run allocations. If firms produce at close to full capacity, contracts signed with one buyer can limit their ability to meet demand shocks elsewhere, while the pro-competitive effects of contracts are likely to be limited, since firms are already producing almost at full capacity. In this case we would expect contracts to hurt short-run allocations.

In Appendix B.1, I carry out Monte Carlo simulations of the spot market model with two buyers and two sellers to illustrate these two mechanisms. I first consider a scenario where firms have large amounts of excess capacity, and show that contracts improve welfare by raising total production. I then consider a scenario where firms are highly capacity-constrained, so that they cannot adjust their total production. I show that in this case, contracts lead to reduced welfare, and the reduction in welfare comes about because contracts lead to distorted allocations during periods with high demand in one market and low demand in another.

Contracting Externalities

A second feature of the model is the presence of externalities. A contract signed between a particular seller and a buyer causes the spot market to become thinner, with both reduced supply and demand. Due to demand uncertainty, this raises the expected spot price paid by every other buyer; thus, contracts impose negative externalities on buyers not party to the contract.

In the bargaining model, each seller-buyer pair chooses the contract quantity that maximizes their bilateral joint surplus, without accounting for the effect of these contracts on other buyers. As a consequence, the equilibrium level of contracting is generically not going to be socially optimal. To see this, consider a situation where there is one seller i and two buyers j and j' , and they sign contracts in Stage 3 (after the seller commits to the investment). The contract quantity q_{ij} between seller i and buyer j satisfies the FOC:

$$\frac{\partial V_i^3}{\partial q_{ij}} + \frac{\partial v_j^3}{\partial q_{ij}} + \frac{\partial \bar{w}_j(q_{ij})}{\partial q_{ij}} + \eta_{ij}^3 = 0 \quad (2.26)$$

However the contract quantity between seller i and buyer j that would maximize the joint surplus of seller i , buyer j and buyer j' would satisfy a different FOC:

$$\frac{\partial V_i^3}{\partial q_{ij}} + \frac{\partial v_j^3}{\partial q_{ij}} + \frac{\partial \bar{w}_j(q_{ij})}{\partial q_{ij}} + \eta_{ij}^3 + \frac{\partial v_{j'}^3}{\partial q_{ij}} = 0 \quad (2.27)$$

Thus as long as the payoff of buyer j' , $v_{j'}^3$, depends on the contract quantity signed by buyer j (q_{ij}), the equilibrium level of contracting will not be socially optimal. If $\frac{\partial v_{j'}^3}{\partial q_{ij}} < 0$, that is if contracts impose negative externalities, then there will be over-contracting in equilibrium.

2.5 Estimation

I now describe how to estimate the model described in Section 2.4 and present results from the estimation. I begin by describing estimates of demand and production costs. I then describe estimation of the contracting and investment model.

2.5.1 Estimation of Demand Curve

I estimate the following demand curve for importing country j :

$$Q_{jt} = -bp_{jt} + \theta_j + x_{jt}\theta_{dx} + \varepsilon_{jt} \quad (2.28)$$

θ_j denotes country fixed effects and x_{jt} denotes demand shifters. In the baseline specification, this includes electricity consumption from fossil fuels.²⁴ Liquefied natural gas is almost always converted back into gaseous form before it is consumed, and natural gas is primarily used for power generation and heating. Electricity consumption thus captures changes in electricity demand that translate into a shift in the demand curve for LNG. I use country fixed effects to capture differences in heating demand across countries as well as other time-invariant country-specific factors that affect LNG demand (for example the amount of piped natural gas the country has access to).

Estimation of equation (2.28) via ordinary-least squares runs into standard endogeneity problems because the spot price p_{jt} is positively correlated with the unobserved demand shock ε_{jt} : the greater country j 's demand for LNG, the higher the equilibrium spot price in country j . I consider two instruments for p_{jt} . The first

²⁴The results are very similar if we use the country's total electricity consumption (from all energy sources) instead.

Table 2.4: Demand Curve Estimates

	(1)	(2)	(3)	(4)
	OLS	IV	IV	IV
Spot Price	0.70*** (0.24)	-1.54 (1.26)	0.13 (0.27)	-0.37*** (0.10)
Electricity Consumption			0.0039* (0.0022)	0.016*** (0.0023)
Importer Fixed Effects	No	No	No	Yes
N	317	317	297	297
R^2	0.026	-0.24	0.060	0.96
Cragg-Donald Wald F-statistic		17.0	32.9	33.4
Kleibergen-Paap F-statistic		13.0	87.0	65.5

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

instrument, which varies only over time but not across j , is total liquefaction capacity (i.e. total export capacity) in the world in period t . The greater LNG capacity in period t , the higher is the supply of LNG and therefore the lower the price in period t . The identification assumption is that LNG export capacity is uncorrelated with idiosyncratic demand shocks today, after controlling for electricity consumption from fossil fuels. The logic behind the instrument is that LNG terminals take many years to build, and at the time the decision to invest is made, it is difficult to foresee idiosyncratic demand shocks that are realized several years later. The modern history of the LNG industry is replete with instances where sellers make investments without fully anticipating how demand would evolve in the importing countries. For example, Qatar's massive capacity expansion in the 2000s was driven to a large degree by the expectation that the US would be a major importer of natural gas. However by the time all of Qatar's terminals came

online, US demand for LNG had shrunk dramatically due to the shale gas boom, and instead Qatar ended up turning to Asian countries as its major buyers of LNG.

The second instrument, which varies over both t and j , is electricity demand in rival countries. The instrument is relevant because if other countries experience higher electricity demand, demand for LNG rises, and the price paid by country j will increase. The instrument is plausibly exogenous provided that country j 's idiosyncratic demand shocks, after controlling for its own electricity demand, are uncorrelated with electricity demand in other countries.

The demand curve estimates are shown in Table 2.4. Column (1) shows the OLS regression of demand on spot prices, without any controls. As expected, the coefficient is positive, reflecting the fact that price changes are mostly driven by shifts in demand rather than shifts in supply. In Column (2)-(4), I instrument for spot prices using total liquefaction capacity and total electricity demand in countries other than country j . The 1st-stage Cragg-Donald Wald F-statistic ranges from 19 to 35, suggesting that weak instruments are unlikely to be a concern. In the most flexible specification (4) where I include both electricity demand and importer fixed effects, the coefficient on the spot price is negative and statistically significant, and implies a demand elasticity of around 0.35 for the average observation. The coefficient on electricity consumption is also positive and statistically different from zero, and implies that the elasticity of LNG demand with respect to electricity consumption is around 0.82 for the average observation.

Table 2.5 shows the effect of adding other demand shifters. In column (2), we add in the Brent oil price: oil and natural gas are substitutes in power generation, so an increase in the oil price increases LNG demand. In column (3), we add in the minimum

annual temperature.²⁵ Natural gas is often used for heating, so that LNG demand spikes during colder winters. Finally, in column (4), we control for both the oil price and the minimum annual temperature. Both of these demand shifters have the expected sign: LNG demand increases if the oil price goes up or if the minimum temperature goes down.

Table 2.5: Demand estimation: adding other demand shifters

	(1) IV	(2) IV	(3) IV	(4) IV
Spot Price	-0.37*** (0.10)	-0.77*** (0.23)	-0.29*** (0.11)	-0.55** (0.26)
Electricity Consumption	0.016*** (0.0023)	0.016*** (0.0025)	0.014*** (0.0027)	0.014*** (0.0029)
Oil Price		0.078*** (0.027)		0.056* (0.033)
Minimum annual temperature			-0.52** (0.20)	-0.38** (0.15)
Importer FE	Yes	Yes	Yes	Yes
N	297	297	251	251
R^2	0.96	0.96	0.96	0.96
Cragg-Donald Wald F-statistic	33.4	19.5	21.8	24.5
Kleibergen-Paap F-statistic	65.5	62.3	44.4	58.2

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

²⁵The temperature data is obtained from the World Bank's Climate Change Knowledge Portal. This does not include weather data for 2017 nor does it include any data for Taiwan, which is why the number of observations decreases from 297 to 251.

2.5.2 Estimation of Cost Parameters

The starting point for the estimation of the cost parameters is the first-order condition (2.7), which I have repeated here:

$$p_{jt}^* + S_{ij} \frac{\partial p_{jt}^*(S_{ijt}, S_{-ijt})}{\partial S_{ijt}} - \frac{\partial C(q_{it}, K_{it})}{\partial S_{ijt}} - (c_{ijt}^d + \phi(x_{ijt})) \leq 0 \quad (2.29)$$

As discussed earlier, I assume that firms face a hockey-stick cost function, where the marginal cost is flat until the firm hits a capacity utilization of v , at which point the marginal cost curve starts to slope upwards:

$$C(q_{it}, K_{it}) = \delta_1 q_{it} + \frac{\delta_2}{2} 1(q_{it} \geq vK_{it})(q_{it} - vK_{it})^2 \quad (2.30)$$

Estimation routine

The parameter vector I have to estimate is $\theta = (\phi(\cdot), \delta_1, \delta_2, v)$. I estimate the model via a minimum distance procedure. For each guess of θ , I solve numerically for the sellers' predicted spot sales to each buyer via a fixed point routine. I minimize the sum of squared deviations between the model predicted spot sales and observed spot sales.

Results

In this section I describe estimates of relatively parsimonious versions of this model. I do not include any covariates in x_{ijt} so that the trade costs between i and j are purely given by the (observed) shipping costs, and set $\delta_1 = 0$. In addition I set $v = 0.90$ because I found it is difficult to identify v in practice. The parsimonious models leads to a reasonably good fit of prices and spot exports, and adding more covariates does not measurably increase the model fit.

Table 2.6 presents cost function estimates. The estimated δ_2 of 95 implies that if a shipyard goes from 90% to 100% capacity utilization, marginal costs increase by \$9.5/MMBtu; in other words, the cost of operating at high capacity utilization is very high.


Table 2.6: Conduct and Cost Parameter Estimates	
	(1)
Hockey-stick costs	
Convexity, δ_2	95.49 (7.66)
Hockey-stick threshold, v	0.90
N	3145
Buyers	40
Fit (R^2)	
Prices p_j	0.63
Total Exports by Exporter q_i	0.96
Total Spot Exports by Exporter $\sum_j S_{ij}$	0.40
Spot Trade Flows S_{ij}	0.35

Note: Each observation is an exporter-importer-year pair. Total LNG spot exports from country i to country j in year t are measured in million tonnes. The spot price and shipping costs are measured in \$/MMBtu. Heteroskedasticity-robust standard errors reported in parentheses.

The model fits many of the key data patterns reasonably well, as shown in the last panel of Table 2.6. The R^2 for prices is 0.63, while the R^2 for spot trade flows is 0.35. The R^2 for total exports is close to 1, but this is because total exports are largely pinned down by contracts. Figure 2.12 shows actual versus predicted prices in Japan. We can see that the model generates a reasonable fit for the year-to-year changes in prices.

In robustness analyses, I have also worked with a quadratic cost function where the marginal cost increases smoothly with the level of capacity utilization (i.e. q/K). I found that the hockey-stick model does a much better job than this alternative model at

Figure 2.12: Actual vs. Predicted Spot Prices in Japan



Note: The figure plots actual spot prices in Japan and as well as model-predicted annual spot prices in Japan.

predicting the sharp price movements from year to year. Under the quadratic model, because the cost function is smooth, production is not as constrained by available capacity as with the hockey-stick model. Under the hockey-stick model, by contrast, because capacity constraints are much more tight, production is relatively inflexible and thus prices move much more, and this appears to better pick up the reality of the industry.

2.5.3 Estimation of contracting and investment model

I now discuss how to estimate the contracting and investment model, or Stages 1 - 3 of the four-stage game. The estimation utilizes the first-order conditions of each of these three stages of the game, as shown below:

$$\frac{\partial V_i^3}{\partial q_{ij}} + \frac{\partial v_j^3}{\partial q_{ij}} + \frac{\partial \bar{w}_j(q_{ij})}{\partial q_{ij}} + \eta_{ij}^3 = 0 \quad (2.31)$$

$$\frac{\partial V_i^2}{\partial K_i} - \frac{\partial \bar{C}_i(K_i)}{\partial K_i} - \eta_i^2 = 0 \quad (2.32)$$

$$\frac{\partial V_i^1}{\partial q_{ij}} + \frac{\partial v_j^1}{\partial q_{ij}} + \frac{\partial \bar{\omega}_j(q_{ij})}{\partial q_{ij}} + \eta_{ij}^1 = 0 \quad (2.33)$$

Equation (2.31) is the first-order condition of the Stage 3 contracting problem, which states that the marginal value of contracting to the buyer and seller should equal 0 at the optimal contract quantity. Likewise equation (2.33) is the first-order condition of the Stage 1 contracting problem. Finally, equation (2.32) is the first-order condition for the seller's investment problem, which states that the marginal benefit to the seller from investing equals the marginal cost of investment.

In order to take the theoretical model to the data, I have to make parametric assumptions about the shape of the investment cost function as well as the contract premium. The particular parametrizations I have adopted are:

$$\begin{aligned} \bar{w}_j(q) &= \theta_3 q \\ \bar{C}_i(K) &= c_1 K + \frac{c_2}{2} K^2 \\ \bar{\omega}_j(q) &= \theta_1 q \end{aligned}$$

θ_1 and θ_3 represent the per-unit contract premium that buyers in Stage 1 and 3 respectively receive from signing long-term contracts (relative to spot trade). We allow the contract premium to differ between Stage 1 and 3 because buyer preferences for contracts could depend on when the contract is signed; for example if the motivation is to lock in supplies in advance, and buyers prefer to lock in supplies earlier rather than latter, then θ_1 might be higher than θ_3 . Alternatively, θ_1 might be lower than θ_3 if buyers prefer to wait before signing long-term contracts.

c_1 and c_2 are investment cost parameters, where c_1 represents the marginal cost of

investment at $K = 0$ while c_2 scales how the marginal cost of investment varies with the level of investment.

Finally, I assume that the bargaining weight equals τ for all sellers. In principle it would be ideal to allow the bargaining weight to differ by seller, but in practice it is difficult to identify (in the absence of data on contract prices).

The structural parameters to be estimated are: θ_3 and θ_1 (contract premium in Stage 3 and Stage 1); c_1 and c_2 (investment cost parameters) and τ (bargaining weights).

Approximation of buyer and seller payoffs

In order to estimate the model, for each parameter guess, we need to be able to solve for the seller and buyer payoff functions at each of the three stages of the game. Recall from equation (2.8) and (2.9) that seller and buyer payoffs at the end of Stage 3 are discounted sums of their per-period payoff functions $\pi_{it}^s(q_t^c, K_t)$ and $\pi_{jt}^b(q_t^c, K_t)$ respectively, where π_{it}^s is seller i 's expected profits in period t and π_{jt}^b is buyer j 's expected consumer surplus in period t .

Because of the complex nature of the spot market equilibrium, however, analytical expressions for $\pi_{it}^s(q_t^c, K_t)$ and $\pi_{jt}^b(q_t^c, K_t)$ do not exist. This means that constructing the payoffs from the three-stage game is computationally very demanding, since it requires solving the spot market equilibrium for every year while integrating out demand shocks.

Instead, I use approximations of the per-period expected payoff functions when estimating the model. 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 $\phi_1, \dots, \phi_{L_b}$. These approximate per-period payoff functions

are shown in the following equations:

$$\pi_{it}^s(q_t^c, K_t) \simeq \sum_{l=1}^{L_s} b_l^s u_l(q_t^c, K_t) \quad (2.34)$$

$$\pi_{jt}^b(q_t^c, K_t) \simeq \sum_{l=1}^{L_b} b_l^b \phi_l(q_t^c, K_t) \quad (2.35)$$

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

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

The choice of basis functions is crucial in ensuring well-behaved approximations. I approximate buyers' expected payoffs as quadratic functions of their total contracted quantity, rival contracted quantity and total capacity in that year, and interact these terms with the buyer's average demand state and the average distance from the buyer to other sellers. I approximate sellers' expected payoffs as quadratic functions of their total contracted quantity, their capacity, rival contracted quantity and rival capacity in that year, and interact these terms with the average demand state in that year and the average distance from the seller to other buyers. I also include the term $\log(\psi K_i - Q_i^c)$ in the approximation of sellers' payoffs to capture the effects of increasing marginal costs as the seller approaches full capacity utilization.

Identification

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

The key identification challenge is that contract prices are unobserved. Identification of τ instead comes from the fact that I observe both investment decisions as well as contract quantity decisions in Stage 1 (before investment is committed to). Sellers in this model are subject to a hold-up problem that is worse the lower their bargaining power. The lower the bargaining power, the lower is the investment they choose in Stage 2, and the higher is the quantity contracted in Stage 1 (when they are not subject to the hold-up problem).

Separately identifying τ and investment costs requires that we not only see investment data, but also data on pre-investment contracting. In the absence of data on pre-investment contracting, it would be challenging to separately identify investment costs and bargaining power. This is because both high investment costs and low bargaining power imply that investment is relatively unresponsive to changes in the value of investing (e.g. from changes in demand over time).

The Stage 1 contracting moments also help identify θ_1 , the contract premium in Stage 1. Separate identification of θ_1 and τ comes from variation in the extent to which sellers face hold-up risk in Stage 3. A large θ_1 leads to an increase in contracted quantity for all sellers in Stage 1, whereas a lower τ leads a larger increase in contracted quantity for sellers that are more exposed to hold-up risk (for example sellers that are located

further away from alternative buyers).

Finally, Stage 3 contract decisions are made after the sellers commit to the investment; moreover, the contract quantity is chosen to maximize the joint surplus of the buyer and seller, irrespective of their bargaining weights. Thus the Stage 3 contract quantity moments in equation (2.31) are not informative about investment costs or the bargaining weight, and instead help to pin down θ_3 or the contract premium.

Results

I estimate the model via non-linear least squares. The NLLS estimator minimizes the weighted sum of squared residuals, where the weights are the inverse of the variance of the moments.

Table 2.7 shows the estimated parameters. In Column (1), we use all the FOCs, including the pre-investment contract quantity decisions. We find a bargaining weight of 0.55. This suggests that buyers on average have sizeable bargaining power, and we can reject the hypothesis that sellers have the ability to make take-it-or-leave-it offers to the buyer. This echoes much of the empirical Nash bargaining literature that generally finds firm behavior is inconsistent with the take-it-or-leave-it bargaining model (Crawford and Yurukoglu, 2012; Grennan, 2013; Ho and Lee, 2017).

We find a θ_3 of around 0.74; this means that on average, buyers are willing to pay a premium of about \$0.74/MMBtu to trade via contracts signed in Stage 3 rather than via spot. The estimated contract premium for Stage 3 is fairly modest in relation to the average price of LNG (typically around \$9/MMBtu). θ_1 is estimated to be negative at around \$-0.06/MMBtu, meaning that for Stage 1 there is little evidence of any contract premium at all. The contract premium appears to be higher for Stage 3 than for Stage 1,

Table 2.7: Parameter estimates using minimum distance estimator

	(1)	(2)
Contract premium post, θ_3	0.74 (0.19)	0.74 (0.19)
Contract premium pre, θ_1	-0.06 (0.48)	
Investment cost, c_1	99.44 (6.69)	110.73 (21.10)
Investment cost, c_2	-0.96 (0.29)	-2.75 (2.23)
Seller bargaining weight, τ	0.55 (0.11)	1.55 (2.13)
Number of contracts (Stage 3)	172	172
Number of contracts (Stage 1)	123	123
Number of investments	54	54

Note: Column (1) reports parameter estimates using all three sets of moments: equations (2.31), (2.32) and (2.33). Column (2) reports parameter estimates using only the moments corresponding to Stage 2 and Stage 3: equations (2.31) and (2.32). All parameters estimated using non-linear least squares. Standard errors are heteroskedasticity-robust.

though the difference between θ_3 and θ_1 is not statistically different from 0.

The marginal cost of investment c_1 is estimated to be around 100 in Column (1). Finally c_2 is estimated to be around -1, meaning the cost function is concave. This suggests there are economies of scale in building LNG projects. I discuss below the interpretation of the investment cost parameters.

In Column (2), I report the results when we estimate the model using only the FOCs of the Stage 2 and Stage 3 decisions, ignoring the information provided by pre-investment contracting in Stage 1. Comparing Column (1) and Column (2), we can see that utilizing the data on pre-investment contracting is important in order to obtain sensible estimates of the parameters. In Column (2), the seller bargaining weight is

now estimated to be implausibly high at 1.55, and is imprecisely estimated.²⁶ The estimated investment cost is also significantly higher. This suggests that in the absence of data on pre-investment contracting, the model is unable to distinguish between low seller bargaining power and high investment costs, since both of these lead to under-investment. It is only by observing the volume of pre-investment contracting that we can identify the bargaining weight separately from investment costs.

Discussion

I now discuss what the estimated model implies for the cost of investment, investment incentives and social optimality of contracting.


Using the NLLS estimate from column (1) as a benchmark, the estimated average cost of investment for a 7 MTPA project (an average-sized project) is \$96/MMBtu. This means that for such a project, the average cost of building 1 MTPA of capacity equals \$4.7 billion. Because the parameter c_2 is negative, there are economies of scale in plant construction. These are fairly modest in magnitude: the largest plant in the data (with size 16 MTPA) is estimated to have an average cost of \$4.5 billion/MTPA, while the smallest plant in the data (with size 0.5 MTPA) is estimated to have an average cost of \$4.9 billion/MTPA.

For 27 of these projects, I have accounting data on the cost of investment. From this data, the average cost of building 1 MTPA of capacity equals \$1.7 billion. So my estimated investment cost is significantly larger than the accounting cost. Figure 2.13 shows how the estimated investment cost differs from the reported accounting cost of building the terminals.

²⁶The theoretical range of the bargaining weight is from 0 to 1.

The difference could be because of a few factors. Firstly, in the spot market model, I do not estimate the intercept of marginal costs, nor do I estimate fixed costs of operating a project (or maintenance). It is likely that these are instead being picked up in the estimated investment cost. In addition, the “economic” cost of investment includes the opportunity cost of using these funds elsewhere. Given that these are massive projects requiring large loans and complex financing deals, the interest payments on these loans are also likely to be quite high. There is considerable risk involved in these projects and this might be getting picked up in the cost estimates.

Figure 2.13: Estimated cost of investment vs. reported cost of investment



The estimated parameters suggest sellers face substantial hold-up risk. In the model, sellers build capacity up to the point where their marginal private benefit from investing (which includes the transfer received from the buyer in Stage 3) equals the marginal cost of investing, as we can see from equation (2.32). Because buyers have substantial bargaining leverage, though, sellers do not receive the full marginal benefit of investing. If buyers 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 20% on average. Thus sellers’ incentives to invest are

significantly dampened by the existence of the hold-up effect.

Finally, we can use the estimated model to assess the direction and magnitude of contracting externalities, building on the discussion on p. 100. For each contract, I compute $\frac{dW}{dq^c}$, or the derivative of social welfare with respect to the contract quantity. If the level of contracting reached in equilibrium were socially optimal, this quantity would equal 0 on average. Instead, I find that $\frac{dW}{dq^c}$ is on average equal to \$1.5/MMBtu, implying that the marginal external cost of contracting is sizeable and equal to about 17% of the average price of LNG on the spot market. This implies that a 1% increase in the contract quantity would decrease aggregate welfare by about 0.1%.

2.6 Counter-factual Analysis

Using the estimated model, I carry out two counter-factual exercises in order to explore the consequences of long-term contracts for efficiency. In the first counter-factual, I hold industry capacity fixed and reduce the quantity traded under every long-term contract by 50%. This allows me to measure the allocative efficiency consequences of using long-term contracts. I find that over the period from 2006 to 2017, aggregate welfare would increase by \$27 bn relative to the benchmark scenario, or around 2% of industry revenue. The welfare gains would be larger (at around \$70 bn) if the spot market were competitive: as pointed out by Allaz and Vila (1993), contracts have a pro-competitive effect when sellers have market power on the spot market, and so reducing the degree of contracting leads to sellers having a stronger incentive to price discriminate on the spot market. However I find that the net benefits from greater inflexibility outweigh the potential losses from increased market power.

In the second counter-factual, I shut down the ability of sellers and buyers to sign

long-term contracts, and solve for each sellers' investment decisions (where they can only sell LNG on the spot market). When implementing this counter-factual, I hold fixed choices of sellers and buyers in other projects. Thus, the counter-factual identifies the partial equilibrium effect of shutting down contract usage.

I find that the average seller reduces investment by 35%, from 6.61 million tonnes per annum to 4.30 million tonnes per annum if they were not allowed to sign long-term contracts with buyers, as we can see from Table 2.8.

Table 2.9 shows that the largest reduction in investment in this counter-factual comes from sellers who are on average located further away from buyers than the median seller: these are the sellers that are most subject to hold-up risk, and as such are the ones that are least likely to invest without the protection afforded by long-term contracts. As we might expect, sellers that sign a large share of their capacity under long-term contracts also experience bigger reductions in capacity in a world where they could sign long-term contracts.

Table 2.8: Average capacity with and without contracting

	Benchmark	No contracting
Capacity of average project (MTPA)	6.61	4.30

Table 2.9: Decomposing capacity reduction by type of seller

	Top 50%	Bottom 50%
Share of capacity contracted	55%	17%
Total capacity (size)	37%	35%
Average distance from buyers	52%	29%

2.7 Conclusion

Long-term contracts are useful because they allow firms to make sunk cost investments even in the presence of holdup risk, but rigidities in contract design can inhibit the ability of firms to respond to demand shocks. In this paper, I develop and estimate a structural model of investment, contracting and spot trade to quantify the trade-off firms in the LNG industry face between hold-up risk and contract rigidity. I show that the timing of contracting and investment decisions can help identify bargaining power even when contract prices are unobserved. My estimates imply that buyers have fairly substantial bargaining leverage, implying that sellers face holdup risk when investing. Reducing the use of long-term contracts would improve short-run allocative efficiency, but lead to a substantial reduction in investment.

My findings suggest that policies that are aimed at improving allocative efficiency may not necessarily improve welfare if they are not accompanied by provisions to safeguard the incentives of agents to make sunk cost investments. In ongoing work, I assess whether one such policy, the removal of destination clauses from long-term LNG contracts, can raise welfare. In a world without destination clauses, the cost of engaging in arbitrage would be substantially lower and sellers' ability to engage in spatial price discrimination would be reduced, which would likely result in improved allocative efficiency. However, a potential side-effect of such a policy would be to reduce the share of profits accruing to sellers and increase the bargaining leverage of buyers, potentially resulting in a reduced incentive to invest. Thus whether or not such a policy is welfare-improving depends on whether the improvement in allocative efficiency dominates any loss in welfare coming from reductions in investment.

CHAPTER 3

CHINA'S INDUSTRIAL POLICY: AN EMPIRICAL EVALUATION

Panle Jia Barwick, Myrto Kalouptsidi, Nahim Bin Zahur

Abstract: Despite the historical prevalence of industrial policy and its current popularity, few empirical studies directly evaluate its welfare consequences. This paper examines an important industrial policy in China in the 2000s, aiming to propel the country's shipbuilding industry to the largest globally. Using comprehensive data on shipyards worldwide and a dynamic model of firm entry, exit, investment, and production, we find that the scale of the policy was massive and boosted China's domestic investment, entry, and world market share dramatically. On the other hand, it created sizable distortions and led to increased industry fragmentation and idleness. The effectiveness of different policy instruments is mixed: production and investment subsidies can be justified by market share considerations, but entry subsidies are wasteful and dissipated by the entry of inefficient firms. Finally, how the policy is implemented matters: the distortions could have been significantly reduced by implementing counter-cyclical policies and by targeting productive firms.

3.1 Introduction

Industrial policy, broadly defined as a policy agenda aimed at shaping a country's industrial structure by promoting certain sectors, has been widely used in developed and developing countries. Examples include the U.S. and Europe after World War II, Japan in the 1950s and 1960s (Johnson, 1982; Ito, 1992), South Korea and Taiwan in the 1960s and 1970s (Amsden, 1989), and China, India, Brazil, and others more recently

(Stiglitz and Lin, 2013). In recent years, industrial policies have reemerged and been openly embraced by policymakers around the world. As Rodrik (2010) puts it, “The real question about industrial policy is not whether it should be practiced, but how.”

Despite the prevalence of industrial policies, we have limited empirical evidence regarding the efficacy of different policy instruments and their impact on industrial evolution. Our paper contributes to the existing literature on industrial policy by focusing on China, currently the most prominent adopter of industrial policies. During the past two decades or so, Chinese firms have rapidly dominated many global industries, such as steel, auto, and solar panels (see Figure C.1), partly as a result of government support.

In this paper, we analyze the shipbuilding industry, an illustrative example of the quick ascent of China’s manufacturing sector into global influence during the 2000s. At the turn of the century, China’s nascent shipbuilding industry accounted for less than 10% of world production. During the 11th (2006-2010) and 12th (2011-2015) National Five-Year Plans, it was dubbed a pillar industry in need of special oversight and support. Since then, a remarkable number of national policies were issued to develop the infant industry into the largest worldwide. Within a few years, China overtook Japan and South Korea and became the leading ship producer in terms of output. At the same time, the industry witnessed a massive entry wave of new firms, which led to a fragmented industry structure. The resulting low capital utilization proved detrimental during the recession and triggered the adoption of consolidation policy by the central government.

Using this historical event as a case study, we address two questions of interest. First, how did China’s industrial policy affect the evolution of the domestic and global industry? Second, what is the relative performance of different policy instruments, which include production subsidies (e.g., subsidized material inputs, export

credits, buyer financing), investment subsidies (e.g., low-interest long-term loans, expedited capital depreciation), entry subsidies (e.g., below-market-rate land prices), and consolidation policies (White Lists)? Given the pronounced firm heterogeneity and salient business cycles in this industry (as in most manufacturing industries), there could be substantial wedges among different policies, and hence room for improvement in terms of the policy design.

A critical challenge in addressing our questions is the lack of suitable data. Although we have access to firm-level data on output quantity, ship price, capital, and firm characteristics, we lack information on the nature and magnitude of government subsidies. This is not specific to our analysis. As documented in a large literature (Bruce, 1990; Young, 2000; Anderson and Van Wincoop, 2004; Poncet, 2005; Kalouptsi, 2018; Barwick *et al.*, 2020; Bai and Liu, 2019), industrial subsidies are often purposefully covert and notoriously difficult to detect and measure. To overcome this challenge, we thus adopt the approach in the existing literature and recover the subsidy magnitude by estimating the cost structure of the industry before and after the policies were implemented.

Industrial policies affect industry evolution and entail long-lasting consequences. To evaluate these long-term implications, we build a dynamic model that centers on firm production, investment, entry, and exit. In each period, firms optimize production to maximize per period variable profits. Then they decide whether or not to exit and how much to invest conditional on staying. Investment increases firm capital, which in turn reduces future production cost. Potential entrants make entry decisions based on their expected lifetime profitability, as well as the cost of entry. China's industrial policy affects all of these decisions by lowering the relevant costs.

We estimate the model primitives and compute the industry evolution under

alternative policy scenarios. The main objects of interest are production costs, investment costs, entry costs, exit scrap values, as well as the demand for new ships. Estimates for industrial subsidies are obtained via changes in firm costs upon the introduction of different policies. Our empirical strategy relies on the literature of industry dynamic models, with a couple of departures. First, we estimate fixed costs of production to accommodate periods with zero production, when idling capacity plagued this industry in the latter part of our sample. Second, we allow for continuous investment under unobserved cost shocks and adjustment costs, in light of heterogeneous investment across firms with similar attributes.¹

Our analysis delivers four sets of main findings. First, like many other policies unleashed by China's central government in the past few decades, the scale of the industrial policy in the shipbuilding industry is massive (relative to the size of the industry). Our estimates suggest that the policy support from 2006 to 2013 is equivalent to RMB 460 billion, with the lion's share going to entry subsidies (RMB 297 billion), followed by production subsidies (RMB 125 billion) and investment subsidies (RMB 37 billion).² It boosted China's domestic investment and entry by 270% and 200%, respectively, and increased its world market share by 40%; importantly, three-quarters of this expansion occurred via business stealing from rival countries. However, the policy generated merely RMB 45 billion of net profit gains to domestic producers and RMB 230 billion of worldwide consumer surplus. The long-term rate of return of the adopted policies, in terms of net lifetime profit gains to domestic producers, is only 16%. Fixed costs are an important contributor to the low return, due to the volatile nature of the industry. The policy attracted a large number of inefficient producers and did not

¹Ryan (2012) is another example with continuous investment and heterogeneous shocks. Unlike our setting, he formulates investment as following an S-s rule.

²All monetary values reported in this paper are discounted and deflated to the 2006 RMB. The conversion rate for this period was 6.88 RMB for 1 U.S. dollar.

translate into significantly higher industry profits over the long run.

Second, the effectiveness of different policy instruments is mixed. Production and investment subsidies can be justified on the grounds of revenue considerations, but entry subsidies are wasteful and lead to increased industry fragmentation and idleness. This is because entry subsidies attract small and inefficient firms; in contrast, production and investment subsidies favor large and efficient firms that benefit from economies of scale. As expected, production subsidies are more effective at achieving output targets, while investment subsidies facilitate more significant capital formation over the long run. In addition, distortions are convex so that the rate of return deteriorates when multiple policies interact.

Third, our analysis suggests that the efficacy of industrial policy is significantly affected by the presence of boom and bust cycles, as well as by heterogeneity in firm efficiency, both of which are notable features of the shipbuilding industry. A counter-cyclical policy would have out-performed the pro-cyclical policy that was adopted by a large margin. Indeed, their effectiveness at raising long-run industry profit differs by nearly threefold, which is primarily driven by two factors: a composition effect (more low-cost firms operate in a bust compared to a boom) and the much cheaper expansion during recessions. In a similar vein, had the government targeted subsidies towards more efficient firms, the policy distortions would have been considerably lower.

Fourth, we examine the consolidation policy adopted in the aftermath of the financial crisis, whereby the government implemented a moratorium on entry and issued a “White List” of firms chosen for government support. This strategy was adopted in several industries to curb excess capacity and create conglomerates that can compete globally.³ Consistent with the evidence discussed above, we find that targeting low-cost firms

³See <https://www.wsj.com/articles/SB10001424127887324624404578257351843112188>.

creates less distortion; that said, the government's White List was sub-optimal and favored state-owned enterprises (SOEs) at the expense of the most efficient firms.

Our base specification assumes a competitive equilibrium outcome with no market failures. As a result, industrial policy is necessarily welfare-reducing and our results speak to its costs. We have explored several traditional rationales for industrial policy. We find limited evidence that the shipbuilding industry generates significant spillovers to the rest of the domestic economy (e.g., steel production, ship owning, or the labor market). Nor do we find evidence of industry-wide learning-by-doing (Marshallian externalities). In a robustness analysis, we relax the assumption of perfect competition and allow firms to engage in Cournot competition. This makes little difference to our estimates of production costs, because the industry is unconcentrated and even the largest firm commands a modest market share. The absence of significant rents associated with market power implies the strategic trade benefits from industrial policy are very limited (Dixit, 1984; Brander, 1995). In terms of the benefits to Chinese trade, the policy implemented in the shipbuilding industry may have lowered freight rates and boosted China's imports and exports, though evaluating the associated welfare gains requires a general equilibrium trade model and falls beyond the scope of this paper. Finally, it is worth noting that non-economic arguments, such as national security, military considerations, and the desire to be world number one (Grossman, 1990), could also be relevant in the design of this policy. Regardless of the motivation, our analysis estimates the policy's costs and assesses the relative efficacy of different policy instruments that can guide future policies.

Literature Review There is a large theoretical literature on industrial policy (Hirschman, 1958; Baldwin, 1969; Krueger, 1990; Krugman, 1991; Harrison and Rodriguez-Clare, 2010; Stiglitz *et al.*, 2013; Liu, 2019; Itskhoki and Moll, 2019).

The earlier empirical literature on industrial policy mostly focuses on describing what happens to the benefiting industries (or countries) with regards to output, revenue, and growth rates (Baldwin and Krugman, 1988; Hansen *et al.*, 2003; Head, 1994; Luzio and Greenstein, 1995; Irwin, 2000), while recent studies recognize the importance of measuring the impact on productivity and cross-sector spillovers (Aghion *et al.*, 2015; Lane, 2017; Liu, 2019). A related literature analyzes trade policies, in particular export subsidies (Das *et al.*, 2007), R&D subsidies (Hall and Van Reenen, 2000; Bloom *et al.*, 2002; Wilson, 2009), place-based policies targeting disadvantaged geographical areas (Neumark and Simpson, 2015; Criscuolo *et al.*, 2019), and environmental subsidies (e.g., renewable energy subsidies, Yi *et al.* 2015; Aldy *et al.* 2018). To the best of our knowledge, our paper provides the first analysis of the impact of a large-scale industrial policy on the global ship-building industry and evaluates different policy instruments using firm-level data.

Our paper contributes to the growing literature studying how China's industrial development has been shaped by a variety of policy interventions, including industrial policies (Aghion *et al.*, 2015; Liu, 2019) and value-added (VAT) tax reforms that stimulated firm investment (Liu and Mao, 2019; Chen *et al.*, 2019a; Bai and Liu, 2019; Chen *et al.*, 2019b). Our study also builds on an emerging literature on the shipbuilding industry (Thompson, 2001, 2007; Hanlon, 2018). Kalouptsidi (2018) is closely related to our paper and detects the Chinese government's *production* subsidies in the Handysize segment, the smallest bulk ships. In addition to firm production, we also examine firm entry, exit, and investment that are abstracted away in the previous study. Moreover, we focus on assessing the performance of various policy instruments and exploring different aspects of policy design, such as implementation over the business cycle and firm targeting. Finally, we study the entire global shipbuilding industry, including all dry bulk ships (not just Handysize), as well as tankers and containerships.

Methodologically, we build on the literature on dynamic estimation, including Bajari *et al.* (2007); Akerberg *et al.* (2007); Pakes *et al.* (2007); Xu (2008); Aw *et al.* (2011); Ryan (2012); Collard-Wexler (2013); Sweeting (2013); Barwick and Pathak (2015); Fowlie *et al.* (2016), as well as the macro literature on firm investment (Abel and Eberly, 1994; Cooper and Haltiwanger, 2006). Complementing the macro literature that focuses on inaction (zero investment) and adjustment costs, our approach can rationalize different investments chosen by observably similar firms, while at the same time accommodates inaction and adjustment costs. Our analysis of firm investment builds on Akerberg *et al.* (2007) and provides one of the first empirical applications of this model with continuous investment. This approach can be used in a variety of settings where heterogeneity in investment is an important consideration.

The rest of the paper is organized as follows. Section 3.2 provides an overview of China's shipbuilding industry and discusses the relevant industrial policy and our datasets. Section 3.3 presents the model. Sections 3.4 and 3.5 describe the empirical strategy and results. Section 3.6 quantifies the impact of China's government support on industry evolution and evaluates the performance of different policy instruments. Section 3.7 concludes.

3.2 Industry Background and Data

3.2.1 Industry Background

Shipbuilding is a classic target of industrial policy, as it is often seen as a strategic industry for both commercial and military purposes. During the late 1800s and early 1900s, Europe was the dominant ship producer (especially the UK). After World War II,

Japan subsidized shipbuilding along with several other industries to rebuild its industrial base and became the world's leader in ship production. South Korea went through the same phase in the 1970s and 1980s. In the 2000s, China followed Japan and South Korea and supported the shipbuilding industry via a broad set of policy instruments.

The scope of national policies issued in China in the 2000s, especially after 2005, to support its domestic shipbuilding industry is unprecedented. In 2002, when former Premier Zhu inspected the China State Shipbuilding Corporation (CSSC), one of the two largest shipbuilding conglomerates in China, he pointed out that “China hopes to become the world's largest shipbuilding country (in terms of output) [...] by 2015.” Soon after, the central government issued the 2003 *National Marine Economic Development Plan* and proposed constructing three shipbuilding bases centered at the Bohai Sea area (Liaoning, Shandong, and Heibei), the East Sea area (Shanghai, Jiangsu, and Zhejiang), and the South Sea area (Guangdong).

The most important initiative was the 11th National Five-Year Economic Plan (2006-2010) which dubbed shipbuilding as a strategic industry. Since then, the shipbuilding industry, together with the marine equipment industry and the ship-repair industry, has received numerous supportive policies. Zhejiang was the first province that identified shipbuilding as a provincial pillar industry. Jiangsu is the close second, and set up dedicated banks to provide shipbuilding companies with favorable financing terms. In the 11th (2006-2010) and 12th (2011-2015) Five-Year plans, shipbuilding was identified as a pillar, or strategic industry by twelve and sixteen provinces, respectively. Besides these Five-Year Plans, the central government issued a series of policy documents with specific production and capacity quotas. For example, as part of the 2006 *Medium and Long Term Development Plan of Shipbuilding Industry*, the government set an annual production goal of 15 million deadweight tonnes (DWT) to be achieved by 2010, and 22

million DWT by 2015. Both goals were met several years in advance. Appendix Table C.1 documents major national policies issued during our sample period.

The government adopted interventions that affected firms along several dimensions. We group policies that supported the Chinese shipbuilding industry into three categories: production, investment, and entry subsidies. Production subsidies lower the cost of producing ships. For instance, the government-buttressed domestic steel industry provides cheap steel, which is an important input for shipbuilding. Besides subsidized input materials, export credits (Collins and Grubb, 2008) and buyer financing in the form of collateral loans provided by local banks constitute other important components of production subsidies.⁴ To help attract buyers, shipyards have traditionally offered loans and various financial services to facilitate purchasing payment. Investment subsidies take the form of low-interest long-term loans and other favorable credit terms that reduce the cost of investment, as well as preferential tax policies that allow for accelerated capital depreciation. Finally, shortened processing time and simplified licensing procedures, as well as heavily subsidized land prices along the coastal regions, greatly lower the cost of entry for potential shipyards.

An often explicitly-stated goal of China's industrial policy is to create large successful firms that can compete against international conglomerates. In the aftermath of the 2008 economic crisis that led to a sharp decline in global ship prices, the government promoted consolidation policies. The *Plan on Adjusting and Revitalizing the Shipbuilding Industry*, implemented in 2009, resulted in an immediate moratorium on entry with increased investment subsidies to existing firms. The most crucial policy for achieving consolidation objectives was the *Shipbuilding Industry Standard and*

⁴Until 2016, the Chinese government provided a range of subsidies for exporters, including reduced corporate income taxes, refund of the value-added-tax, etc. Shipbuilding companies benefit from export subsidies since most of their products are traded internationally.

Conditions (2013), which instructed the government to periodically announce a list of selected firms that “meet the industry standard” and thus receive priority in subsidies and bank financing.⁵ The so called “White List” included sixty firms in 2014 upon announcement.

In this paper, we focus on the production of three ship types: dry bulk, tankers, and containerships, which account for more than 90% of world orders in tons in our sample period. Dry bulk ships transport homogeneous and unpacked goods, such as iron ore, grain, coal, steel, etc., for individual shippers on non-scheduled routes. Tankers carry chemicals, crude oil, and other oil products. Containerships carry containerized cargos from different shippers in regular port-to-port itineraries. As these types of ships carry entirely different commodities, they are not substitutable; we thus treat them as operating in separate markets.

Shipbuilding worldwide is concentrated in China, Japan, and South Korea, which collectively account for over 90% of the world production. We limit our empirical analysis to shipyards in these three countries.

3.2.2 Data

Our empirical analysis draws on a number of datasets. The first dataset comes from Clarksons and contains quarterly information on all shipyards worldwide that produce ships for ocean transport between the first quarter of 1998 and the first quarter of 2014. We observe each yard’s orders, deliveries, and backlog (which are undelivered orders that are under construction) measured in Compensated Gross Tonnage (CGT), for all of

⁵In practice, favorable financing terms and capital market access are often limited to firms on the White List post 2014.

the major ship types, including bulk, tankers, and containers. CGT, which is a widely used measure of size in the industry, takes into consideration production complexities of different ships and is comparable across types. The entry year for a shipyard is defined as the first year it takes an order or the first year it delivers minus two years to account for the time it takes to build a ship, whichever is earlier.⁶

The second data source is the annual database compiled by the National Bureau of Statistics (NBS) on Chinese manufacturing firms with annual sales above five million RMB. For each shipyard and year, we observe its location (province and city) and ownership status (state-owned enterprises (SOEs), privately owned, or joint ventures with foreign investors). We differentiate SOEs that are part of China State Shipbuilding Corporation (CSSC) and China Shipbuilding Industry Corporation (CSIC), the two largest shipbuilding conglomerates in China, from other SOEs. We link firms over time and construct their real capital stock and investment following the procedure described by Brandt *et al.* (2012). We also observe the annual accounting operation costs for each shipyard. One limitation of the NBS database is that data for 2010 are missing. This prevents us from constructing the firm-level investment in either 2009 or 2010, since investment is imputed from changes in the capital stock.

In addition to these firm-level variables, we collect a number of aggregate variables

⁶As an additional measure of firm entry, we successfully extracted the registration information (date and business scope) for 90% of Chinese firms, from the Trade and Industry Bureau database. The registration data is an imperfect measure of entry since firms often register with a wide business scope (ship building and repairs, marine equipment, marine engineering, etc.), and it is difficult to identify firms whose core business is shipbuilding from the registration data alone. In addition, some firms switch from ship repairs and marine equipment to shipbuilding years after registering with the Trade and Industry Bureau. As such, we use the entry year from the Clarkson's database in our main analysis. Nonetheless, the overall entry pattern is similar across these two measures: entry peaked in 2005-2007 and became minimal post 2009.

for the shipbuilding industry, including quarterly global prices per CGT for each of the three ship types.⁷ The steel ship plate price serves as a proxy for changes in the production cost, as steel is a major input in shipbuilding. We merge all datasets to obtain a quarterly panel of Chinese, South Korean, and Japanese shipyards ranging from 1998 to 2013.


3.2.3 Descriptive Evidence and Summary Statistics

Similar to many other manufacturing industries in China, the shipbuilding industry experienced exponential growth since the mid 2000s. Indeed, China became the largest shipbuilding country in terms of deadweight tons in 2009, overtaking South Korea and Japan. Figure 3.1 plots China's rapid ascent into global influence from 1998 to 2013. At the same time, a massive entry wave of new shipyards occurred along China's coastal area. Figure 3.2 plots the total number of new shipyards by year for China, Japan, and South Korea. The number of entrants is modest for Japan (1.4 per year) and South Korea (1.2 per year), partly due to a lack of greenfield sites to build new shipyards. In contrast, the number of new shipyards in China registered a historic record and exceeded 30 per year during the boom years when the entry subsidy was in place. Entry dropped to 15 in 2009 and became minimal within a couple of years of the implementation of the 2009 entry moratorium, as part of the *Plan on Adjusting and Revitalizing the Shipbuilding Industry*.⁸

⁷We experiment with two price indices, real RMB/CGT vs. USD/CGT, and obtain nearly identical results, suggesting that exchange rate fluctuations are not first-order in our analysis.


⁸No new applications were processed post 2009, but projects already approved were allowed to be completed. In addition, firms registered prior to 2009 but engaged in repairs and marine engineering could 'enter' and produce ships post 2009. Both account for the entry (though at a far reduced rate) that we see in Clarksons from 2009 onward.

Figure 3.1: China's Market Share Expansion




Source: Clarkson Research. Market shares computed from total quarterly ship orders.

Figure 3.2: Entry of New Shipyards




Source: Clarksons Research. Number of new shipyards.

Figure 3.3: Investment by Chinese Shipyards



Source: NBS. Total quarterly investment.

Figure 3.4: Ship Prices



Source: Clarksons Research. Average price in USD/CGT.

The rise in entry was accompanied by a large and unprecedented increase in capital expansion (Figure 3.3). The year of 2006 alone witnessed a steep four-fold increase in investment. The capital expansion was universal across both entrants and incumbents and among firms with different ownership status. Indeed, entrants account for 43% of the aggregate investment from 2006 to 2011, with the remaining 57% implemented by incumbents. Private firms, joint ventures, and SOEs account for 25%, 36%, and 38% of total investment, respectively. In addition, the capital expansion was spread out across provinces, though Jiangsu accounted for a disproportionate share of 40% of the

aggregate investment between 2006 and 2011.

The rapid rise in China's production, entry, and investment coincided with the introduction of China's industrial policy for the shipbuilding industry. The global shipbuilding industry went through a boom in the mid-2000s, roughly concurrent with China's initial expansion. As Figure 3.4 shows, ship prices began rising around 2003 and peaked in 2008, before collapsing in the aftermath of the financial crisis and remaining stagnant from 2009 through to 2013. China's production and investment, on the other hand, continued to expand well after the financial crisis.

Table 3.1 contains summary statistics on key variables of interest. There are a large number of firms, with 266 Chinese shipyards, 108 Japanese shipyards, and 46 Korean shipyards. Industry concentration is low, with a world HHI that varies from 230 to 720 during the sample period.

Table 3.1: Summary Statistics

Variable	Obs	Mean	S.D.	Min	Max
All Observations (including zero orders)					
Bulk orders (1000 CGT)	10,101	17.1	51.9	0.0	968.2
Tanker orders (1000 CGT)	10,583	9.6	46.2	0.0	1119.0
Container orders (1000 CGT)	4,813	18.9	93.9	0.0	1644.1
Observations With Positive Orders					
Bulk orders (1000 CGT)	2,316	74.6	86.5	3.9	968.2
Tanker orders (1000 CGT)	1,436	70.4	107.1	0.05	1,119.0
Container orders (1000 CGT)	625	145.3	222.7	2.3	1,644.1
Other Variables					
Bulk backlog (1000 CGT)	10,101	171.4	329.3	0.0	2830.5
Tanker backlog (1000 CGT)	10,583	98.5	315.1	0.0	3840.8
Container backlog (1000 CGT)	4,813	206.6	670.5	0.0	7362.8
Investment (mill RMB)	4,386	18.5	88.9	-240.5	1,770.7
Capital (mill RMB)	6,157	392.0	806.9	0.3	8,203.3

Note: Summary statistics for shipyards in China, Japan, and South Korea. Investment and capital are limited to Chinese yards.

An important feature of ship production is that shipyards take new orders

infrequently, about 23% of the quarters for bulk and less frequently for tanker and containerships. From 2009 onwards, during a prolonged period of low ship prices, the frequency with which yards took new orders was significantly lower. This lumpiness in ship orders, combined with Chinese shipyards becoming increasingly vulnerable to long periods of inaction in the latter part of our sample, is a key feature of the shipbuilding industry that informs our modeling choices in Section 3.3.

Finally, about 52% of firms in our sample produce one ship type, a pattern that holds across countries. The fraction of ships that produce all three ship types is higher in South Korea (28%) and Japan (16%) and lower in China (14%). If a shipyard never takes orders for a certain ship type throughout our sample, it is assumed not to produce this ship type.

3.3 Model

In this section, we introduce a dynamic model of firm entry, exit, and capital investment. In each period, incumbent firms make static production decisions to maximize their variable profit taking global ship prices as given. Then they choose whether or not to exit, and conditional on staying, how much to invest. A pool of potential entrants make one-shot entry decisions based on their expected discounted stream of profits, as well as the cost of entry. At the end of the period, entry, exit, and investment decisions are implemented and the state evolves to the next period.

Time is discrete and is a quarter. In period t , there are $j = 1, \dots, J_t$ firms in the world that produce ships. There are $m = 1, \dots, M$ types of ships, such as dry bulkers, tankers, and containerships. Ships within a type are homogeneous.

Ship Demand In each period, the aggregate inverse demand for ships of type m is given by the function,

$$P_{mt} = P(Q_{mt}, d_{mt}) \quad (3.1)$$

for $m = 1, \dots, M$, where P_{mt} is the market price of ship type m in period t , Q_{mt} is the total tonnage of type m demanded, and d_{mt} are demand shifters, such as freight rates and aggregate indicators of economic activity.

Ship Production Firm j produces q_{jmt} tons at cost:

$$C(q_{jmt}, s_{jmt}, \omega_{jmt}) = c_0 + c_m(q_{jmt}, s_{jmt}, \omega_{jmt})$$

where c_0 is a fixed cost that is incurred even when shipyards have zero production. Fixed costs are often abstracted away in empirical studies, but in later periods of our sample when the aggregate demand for new ships plummeted and many shipyards reported prolonged periods with zero production, fixed costs constitute a substantial fraction of overall costs and should thus not be omitted. They capture wages and compensation for managers, capital maintenance, land usage, etc.

The second term, $c_m(q_{jmt}, s_{jmt}, \omega_{jmt})$, is the standard production cost. We use s_{jmt} to denote firm characteristics (e.g. capital, backlog, age, location, ownership status), as well as aggregate cost shifters that affect all shipyards (e.g. government subsidies, steel prices). In addition, production costs depend on a shock ω_{jmt} : the larger ω_{jmt} is, the less productive the firm is.

Firms are price takers and choose how many tons to produce for ship type m , q_{jmt} , to maximize their profits:

$$\max_{q_{jmt}} \pi(s_{jmt}, \omega_{jmt}) = P_{mt} q_{jmt} - C(q_{jmt}, s_{jmt}, \omega_{jmt})$$

If the optimal production tonnage for type m , q_{jmt}^* , is positive, it satisfies the following first order condition, where $MC_m(q_{jmt}, s_{jmt}, \omega_{jmt})$ is the marginal cost of producing one tonne of type m :

$$P_{mt} = MC_m(q_{jmt}^*, s_{jmt}, \omega_{jmt}) \quad (3.2)$$

Let $s_{jt} = \{s_{j1t}, \dots, s_{jMt}\}$ and $\omega_{jt} = \{\omega_{j1t}, \dots, \omega_{jMt}\}$ denote the union across ship types of observed state variables and unobserved cost shocks, respectively. A firm's total expected profit from all types, before the cost shocks are realized, is given by:

$$\pi(s_{jt}) = E_{\omega_{jt}} \sum_{m=1}^M \pi(s_{jmt}, \omega_{jmt})$$

Note that with a slight abuse of notation, here we use s_{jmt} to denote all observed payoff relevant variables, which include ship prices in addition to cost shifters.

Finally, in each period, the prevailing ship price, P_{mt} , equates aggregate demand and supply, where the aggregate supply is the sum of q_{jmt}^* defined in (3.2).

Investment and Exit Once firms make their optimal production choice, they observe a scrap or sell-off value, ϕ_{jt} , that is distributed i.i.d. with distribution F_ϕ and decide whether to remain in operation or exit. If a firm chooses to exit, it receives the scrap value. If it remains active, it observes a firm-specific random investment cost shock, v_{jt} , that is distributed i.i.d. with distribution F_v and chooses investment i_{jt} at cost $C^i(i_{jt}, v_{jt})$. The amount invested i_{jt} is added to the firm's capital stock next period, which in turn affects its future production costs.

The value function for incumbent firm j is:

$$V(s_{jt}, \phi_{jt}) = \pi(s_{jt}) + \max \left\{ \phi_{jt}, E_{v_{jt}} \left(\max_i \left(-C^i(i, v_{jt}) + \beta E[V(s_{jt+1}) | s_{jt}, i] \right) \right) \right\} \quad (3.3)$$

$$\begin{aligned} &= \pi(s_{jt}) + \max \{ \phi_{jt}, CV(s_{jt}) \} \\ CV(s_{jt}) &\equiv E_{v_{jt}} \left(-C^i(i^*, v_{jt}) + \beta E[V(s_{jt+1}) | s_{jt}, i^*] \right) \end{aligned} \quad (3.4)$$

where $CV(s_{jt})$ denotes the continuation value, which includes the expected cost of optimal investment and the discounted future stream of profits. Crucially, $E_{v_{jt}}$ is the expectation with respect to the random investment cost shock v_{jt} and i^* denotes the optimal investment policy $i^* = i^*(s_{jt}, v_{jt})$.

The optimal exit policy is of the threshold form: the firm exits the market if the drawn scrap value ϕ_{jt} is higher than its continuation value $CV(s_{jt})$. Since the scrap value is random, the firm exits with probability, $p^x(s_{jt})$, defined by,

$$p^x(s_{jt}) \equiv \Pr(\phi_{jt} > CV(s_{jt})) = 1 - F_\phi(CV(s_{jt})) \quad (3.5)$$

where F_ϕ is the distribution of ϕ_{jt} .

Conditional on staying, firm j observes its investment shock, v_{jt} . Its optimal investment $i^* = i^*(s_{jt}, v_{jt})$, which is non-negative, satisfies the first-order condition:

$$\beta \frac{\partial E[V(s_{jt+1}) | s_{jt}, i^*]}{\partial i} \leq \frac{\partial C^i(i^*, v_{jt})}{\partial i} \quad (3.6)$$

with equality if and only if the optimal investment is strictly positive, $i^*(s_{jt}, v_{jt}) > 0$. When the investment costs are prohibitively high or the expected benefit too low, firm j opts for no investment. Capital depreciates at a rate δ that is common to all firms.

We assume that the cost of investment, $C^i(i_{jt}, v_{jt})$ has the following form:

$$C^i(i_{jt}, v_{jt}) = c_1 i_{jt} + c_2 i_{jt}^2 + c_3 v_{jt} i_{jt} + c_4 T_t i_{jt} \quad (3.7)$$

This (quadratic) specification borrows from the macro literature on investment costs (Cooper and Haltiwanger, 2006) with two important differences. First, investment costs depend on the unobserved marginal cost shock v_{jt} . Much of the existing literature has focused on the lumpy nature of investment (inaction) and adjustment costs, but has not modeled heterogeneous investment decisions among observationally similar firms.⁹ Here, we tackle this heterogeneity by introducing v_{jt} that shifts the marginal cost of investment across firms. Note that v_{jt} can also explain inaction: firms with unfavorably large v_{jt} will choose not to invest. In practice, there are many factors that influence firms' investment decisions. Some firms have political connections that grant them favorable access to the capital market (Magnolfi and Roncoroni, 2018), and others might be experienced at sourcing from equipment suppliers at low costs. In our estimation, once we control for v_{jt} , additional adjustment costs, such as $\frac{i_{jt}^2}{k_{jt}}$ and/or a (random) fixed cost, contribute little to the model fit.¹⁰ A second difference from the literature is that we allow government policies T_t to directly affect the marginal cost of investment.

Entry In each period t , $\bar{N}$ potential entrants observe the payoff relevant state variables and their own entry cost κ_{jt} , which is i.i.d., and make a one-time entry decision. The entry cost is drawn from a distribution F_κ that is shifted by the government policy. If potential entrant j decides not to enter, it vanishes with a payoff of zero.¹¹ If j enters, it

⁹Notable exceptions include Ryan (2012) that models firm investment decisions as following an S-s rule and Collard-Wexler (2013) that analyzes discrete investments.

¹⁰The estimated fixed cost of investment, once included, is economically small. Fixed costs are associated with an inaction region where firms will not make investments smaller than a threshold. The larger the fixed cost, the larger the inaction region. In our data, firms do make small investments, which is inconsistent with a large fixed cost.

¹¹Here we follow the bulk of the empirical literature on industry dynamics (Ericson and Pakes, 1995), where the entry decision involves a simple comparison between the value from entering the market and the random entry cost.

pays the entry cost and continues as an incumbent next period. In addition, the entrant is assumed to be endowed with a random initial capital stock that is realized the following period once the firm becomes an incumbent and begins operation.

Potential entrant j solves,

$$\max \{0, -\kappa_{jt} + E [-C^i(k_{jt+1}) + \beta E [V(s_{jt+1})|s_{jt}]] \}$$

where κ_{jt} is the entry cost, k_{jt+1} is entrant j 's initial capital stock in period $t + 1$ after paying a cost of $C^i(k_{jt+1})$. The expectation is taken over entrant j 's information set at time t , which includes all aggregate state variables.

Similar to the exit decision, the optimal entry policy is of the optimal threshold form: a potential entrant enters the market if the entry cost κ_{jt} drawn is lower than the value of entering, i.e.

$$\kappa_{jt} \leq VE(s_{jt}) \equiv E [-C^i(k_{jt+1}) + \beta E [V(s_{jt+1})|s_{jt}]]$$

Since κ_{jt} is random, the potential entrant enters with probability, p_t^e , defined by,

$$p_t^e \equiv \Pr(\kappa_{jt} \leq VE(s_{jt})) = F_\kappa(VE(s_{jt})) \quad (3.8)$$

Industrial Policy Industrial policies affect the costs of production, investment, and entry and are thus part of the payoff relevant state variables, s_{jt} . We assume that these policies are unexpected and perceived as permanent by all shipyards once they are in place. This is consistent with the empirical patterns documented in Section 3.2.3, where the spike in entry and investment coincides remarkably with the timing of these policies.

Equilibrium The Markov Perfect Equilibrium (MPE) of this model is defined as follows:

Definition. An equilibrium of this model consists of policies, $\{q_{jmt}\}_{m=1}^M$, $i^*(s_{jt}, v_{jt})$, $p^x(s_{jt})$, p_t^e , value function $V(s_{jt})$ and prices P_{mt} , such that the production quantity satisfies (3.2) and maximizes the period profit, the investment policy satisfies (3.6), the exit policy satisfies (3.5), the entry policy satisfies (3.8), and ship prices clear the market each period so that aggregate demand equals total supply. Moreover, the incumbent's value function satisfies (3.3) and firms employ the above policies to form expectations.

The key assumption in invoking the MPE concept is that the transition processes of all payoff relevant state variables (including ship prices) satisfy the Markovian property pre and post the policy intervention. Consequently, the equilibrium is stationary conditioning on our state variables (which include dummy variables for different policies) and the value functions are not indexed by t . While China's market share increased substantially during our sample period, such an expansion can be explained by changes in state variables, including demand and cost factors, as well as policy interventions.¹²

Discussion We close this section with a brief discussion on our assumptions. We assume that shipyards are price-takers. This assumption is motivated by the large number of firms in the industry. As discussed above, there are more than four hundred shipyards globally and the market share of the largest firm is less than 5%. The market is thus sufficiently unconcentrated that market power is not a first order consideration. Nonetheless, we consider a variation of our model that incorporates market power with firms being Cournot competitors in the product market. In that case, the optimal production of firm j in period t (when positive) satisfies a variant of the first order condition (3.2) that includes a markup equal to $-q_{jmt} [\partial P(Q_{mt}, d_{mt}) / \partial q_{jmt}]$.

¹²Stationarity does not preclude business cycles, which are typical in the world shipbuilding industry.

We use ship prices to capture the aggregate business cycles. Relaxing the stationarity assumption is a difficult but important topic for future research.

Empirically, this term is small and the estimates are robust when replicated under the Cournot assumption (we discuss this further in Section 3.5). It is also worth noting that on the ship buyer side (shipowners), monopsony power is not a first-order issue as the concentration among ship ownership is also low.¹³

Ships are assumed homogeneous within a type conditioning on size. In prior work (Kalouptsi (2014), Brancaccio *et al.* (2018)), this assumption is substantiated for dry bulkers. These papers show that both in secondary markets for ships, as well as in transport contracts, the majority of price variation is accounted for by a ship's age and aggregate indicators of economic activity. To further substantiate this assumption, we explore a small sample of new ship purchase contracts with detailed price information and ship attributes. Ship type, ship size (measured in Compensated Gross Tonnage), and quarter dummies explain most of the price variation: the R^2 of a hedonic price regression when these are the only regressors is 0.93 for bulk, 0.94 for tankers, and 0.75 for containers. Ship and shipyard characteristics (age, country, number of docks and berths, etc.) have limited explanatory power: adding shipyard fixed effects to the hedonic regressions adds little to the fit with the exception of containers where the R^2 does increase.

We assume away dynamic considerations in production. In practice, producing a ship takes time and the production decision is in general dynamic: production today affects the backlog tomorrow, which affects tomorrow's operation costs and therefore production decisions. However, as documented in Kalouptsi (2018), cost function estimates under static and dynamic assumptions are fairly similar, especially the estimates that reflect the impact of policy interventions on firms' production costs. This is partly because the amount of drastic production expansion seen in practice

¹³The containership segment may be concentrated among the operators, but they lease containerships from many shipowning firms.

cannot be entirely explained by the inter-temporal considerations that arise with a dynamic production. Note that we allow backlogs to directly affect the marginal cost of production, which proxies for the dynamic considerations in production decisions in a reduced-form manner.

Moreover, we model shipyards' production decision in tonnage, rather than number of ships. Modeling the optimization decisions in both the number and size (tonnage) of ships is substantially more complicated. Since our main focus here is entry, exit, and investment, we make the simplifying assumption that shipyards choose production in tonnage to maximize their static profits. This keeps the model tractable. It is worth noting that our estimated production subsidies are consistent with those in Kalouptsi (2018) that models the choice of the number of ships and focuses on shipyards producing a specific size category of dry bulkers, Handysize vessels.

We assume that the cost shocks ω_{jmt} are i.i.d. There are several reasons for this choice. First, ω_{jmt} is estimated to be only moderately persistent, with a serial autocorrelation of 0.28 for bulk, 0.27 for tankers, and 0.39 for containerships. It is worth noting that our estimation procedure on production cost parameters accommodates persistence in ω_{jmt} (see Section 3.4.1). Second, while it is straightforward to estimate the persistence of these shocks using observed quantity choices, incorporating a persistent time-varying unobserved state variable in a dynamic model raises considerable modeling and estimation challenges. For the same reason, the investment cost shocks v_{jt} are assumed i.i.d. Given that aggregate investment increased by more than four-fold within a year upon the announcement of the 11th National Five-Year Plan (Figure 3.3) and that all firms expanded regardless of their efficiency level, firm-specific persistent investment shocks are unlikely a first-order contributing factor to the boom of the capital expansion observed in our sample.

Lastly, the government policies are perceived as permanent by all firms. In contrast, China's economy is experiencing a drastic transformation with a highly dynamic policy environment. This assumption, however, is standard in the literature (Ryan, 2012). Relaxing it and estimating firms' expectations and adaptation to a changing environment is a difficult but important topic for future research (Doraszelski *et al.*, 2018; Jeon, 2018). One (imperfect) approach to proxy for a dynamic policy environment is to use lower discount rates so that future profits are less relevant for today's decisions. We test for the robustness of our results with different discount rates.

3.4 Estimation Strategy

In this section, we present the empirical approach undertaken to uncover the model parameters. The key primitives of interest are: the world demand function for new ships, the shipyard production cost function, the investment cost function, the distribution of scrap values, and the distribution of entry costs. We estimate the heterogeneous production cost function for shipyards in all countries, but only analyze the dynamic decisions (entry, exit, and investment) for Chinese shipyards as there is little entry, exit, and capacity expansion in Japan and South Korea (OECD, 2015, 2016).

We first estimate the demand curve for new ships, as well as the shipyard production costs. We refer to these parameters as the "static parameters". We use these estimates to construct firm profits in each period. Next, we estimate the parameters governing investment costs and the scrap values for exiting firms, i.e. the "dynamic parameters". This constitutes the bulk of our dynamic estimation where we adopt a two-step procedure in the tradition of Hotz and Miller (1993) and Bajari *et al.* (2007). Finally,

we estimate entry costs.¹⁴ We next describe our approach in detail.

This section is self-contained and the reader may omit it and proceed to the results section if desired. Section 3.4.1 discusses estimation of the static parameters (demand and production costs). Sections 3.4.2 and 3.4.2 present the first and second stage of estimating the dynamic parameters (investment cost, scrap values, as well as entry costs).

3.4.1 Estimation of Static Parameters

Demand The demand curve (3.1) for ship type m is parameterized as follows:

$$Q_{mt} = \alpha_{0m} + \alpha_{pm}P_{mt} + d'_{mt}\alpha_{dm} + \varepsilon_{mt}^d \quad (3.9)$$

The demand shifters d_{mt} include freight rates, the total backlog of type m , and some other type-specific variables. Demand for new ships is higher when demand for shipping services is high, reflected in higher freight rates.¹⁵ Conversely, a large backlog implies that more ships will be delivered in the near future, which reduces demand for new ships today. We also control for aggregate indicators of economic activity relevant for each ship type we consider: the wheat price and Chinese iron ore imports for bulk carriers; Middle Eastern refinery production for oil tankers; and world car trade for containerships. In some specifications, we allow for time trends as well. Finally, we allow the price elasticity to change before and after 2006, the main policy year.

Prices are instrumented by steel prices and steel production.¹⁶ Steel is a major input

¹⁴Combining step two and three and estimating all dynamic parameters jointly is more efficient but computationally more challenging.

¹⁵The freight rate measures are the Baltic Exchange Freight Index for bulk shipping, the Baltic Exchange Clean Tanker Index for tankers, and the Containership Timecharter Rate Index for containerships.

¹⁶Other potential instruments include the aggregate number of shipyards, J_t , and the aggregate capital

in shipbuilding and contributes to 13% of the costs (Stopford, 2009). The identification assumption is that steel prices and steel production are uncorrelated with new ship demand shocks ϵ_{mt}^d . This is a plausible assumption because only a modest portion of global steel production is used in shipbuilding and an increase in ship demand ($\epsilon_{mt}^d > 0$) is unlikely to have much impact on steel prices.¹⁷

As there is a single global market for each ship type, the demand curves are estimated from time series variation. To improve the precision of parameter estimates, we restrict some parameters to be the same across ship types and estimate equation (3.9) jointly across the three types using GMM.

Production Cost We parameterize the marginal cost function for type m , $MC_m(q_{jmt}, s_{jmt}, \omega_{jmt})$, as follows:

$$MC_m(q_{jmt}, s_{jmt}, \omega_{jmt}) = \beta_{0m} + s_{jmt}\beta_{sm} + \beta_{qm}q_{jmt} + \omega_{jmt} \quad (3.10)$$

where q_{jmt} denotes tons of ship type m chosen by firm j in period t . It is worth noting that because of time to build, there is a difference between the orders placed in a period, the deliveries, and the production. We use orders as a measure of q_{jmt} , following Kalouptsi (2018). We do so because the number of tons ordered is the relevant quantity decision made by the firm. In addition, our data source reports orders and deliveries instead of production and it is not straightforward to infer production from orders. Last but not least, the decision on orders corresponds to observed prices, while any constructed measure of production does not.

stock. These cost-side instruments shift the industry supply curve and are determined in period $t - 1$, before demand shocks in period t are realized. Results are robust with or without these additional IVs.

¹⁷In addition, internationally traded steel accounts for less than 8% of the volume of goods transported by dry bulkers (UNCTAD, 2018). Thus, changes in the steel price that affect the amount of steel transported by sea are unlikely to directly affect demand for dry bulkers.

State variables s_{jmt} include firm j 's capital and backlog of all ship types. The capital stock, which is controlled by firm j through its investment decisions over time, reduces production costs via the economies of scale. The backlogs also capture economies of scale, as well as learning by doing, and possibly capacity constraints. The vector s_{jmt} in addition contains the age and ownership status, nationality and region (for Chinese firms), a dummy for large firms, the steel price, as well as polynomial terms of these state variables.¹⁸ Finally, s_{jmt} includes dummies for the policy intervention between 2006 and 2008 and then from 2009 onwards. The production cost shock ω_{jmt} is assumed to be normally distributed with mean zero and variance $\sigma_{\omega m}^2$.

The optimal production q_{jmt}^* , if positive, satisfies the first order condition (3.2). Conditioning on observed shipyard attributes, unfavorable cost shocks could lead to zero production. The parameters characterizing shipyards' production costs are $\theta^q \equiv \{\beta_{0m}, \beta_{sm}, \beta_{qm}, \sigma_{\omega m}\}_{m=1}^M$. We estimate these parameters via a Tobit model:

$$L = \prod_{m=1}^M \prod_{q_{jmt}=0} Pr(q_{jmt} = 0 | s_{jmt}; \theta^q) \prod_{q_{jmt}>0} f_q(q_{jmt} | s_{jmt}; \theta^q)$$

It is worth noting that θ^q is consistently estimated even when ω_{jmt} is correlated over time, despite the fact that this likelihood function assumes (erroneously) that ω_{jmt} is i.i.d. (Robinson, 1982).¹⁹ To obtain the standard errors allowing for autocorrelation in ω_{jmt} , we use 500 block bootstraps.

Finally, it should be noted that a firm's production decisions provide no information on its fixed cost c_0 , since the firm incurs this cost regardless of whether it produces.

¹⁸Large firms are defined as the top firms that account for 90% of the aggregate industry revenue from ship production throughout our sample period. There are fifty-five large Chinese shipyards. Adding this variable (on top of capital and other firm attributes) helps to capture unobserved differences across firms, like management skills, political connections, etc. and improves the fit of our model.

¹⁹Intuitively this is similar to how the OLS estimator in the standard linear regression model continues to be consistent (though not efficient) when the errors are non i.i.d.

However, unlike most empirical studies where fixed costs are assumed away, we take advantage of accounting cost data to calibrate c_0 , exploiting the fact that firms report costs incurred even during periods when the production facility is idle. Details on this calibration procedure are reported in Appendix C.1.2. Restricting the fixed cost to zero may bias the counterfactual analyses (Aguirregabiria and Suzuki, 2014; Kalouptsidei *et al.*, 2018); we discuss this issue further in Section 3.6.1.

3.4.2 Estimation of Dynamic Parameters

We now turn to the estimation of the dynamic parameters, i.e. the investment cost, the scrap value distribution, and the entry cost distribution. To estimate these parameters, we rely on firm choices regarding investment, entry and exit. An important complication in doing so is that optimal choices depend on the value function (as well as an unobserved shock in the case of investment), which is unknown. To tackle this challenge, we follow the tradition of Hotz and Miller (1993) and Bajari *et al.* (2007) (henceforth BBL) and estimate the parameters in two stages. In the first stage, we flexibly estimate investment and exit policy functions, as well as the transition process of state variables from the data. Then, we use these estimates to obtain a flexible approximation of the value function. Here we approximate the value function by a set of B-spline basis functions of the state variables, following Sweeting (2013) and Barwick and Pathak (2015). In the second stage, once we have an estimate of the value function, we formulate the likelihood of the observed investment and exit and recover the dynamic parameters of interest. We next describe each step. Appendix C.1 contains additional details.

First Stage

Exit Policy Function Estimating the exit policy function can be done via a number of different approaches (linear probability models, logit or probit, local polynomial regressions, etc.). Here, we perform a probit regression, though results appear robust across different specifications:

$$Pr(\chi_{jt} = 1 | s_{jt}) = \Phi(h(s_{jt}))$$

where χ_{jt} equals 1 if firm j exits in period t , $h(s_{jt})$ is a flexible polynomial of the states, and Φ is the normal distribution. We denote the first-stage estimate of the exit probability by $\hat{p}^x(s_{jt})$.

Investment Policy Function Recall that the cost of investment, $C^i(i_{jt}, v_{jt})$, has the following form:

$$C^i(i_{jt}, v_{jt}) = c_1 i_{jt} + c_2 i_{jt}^2 + c_3 v_{jt} i_{jt} + c_4 T_t i_{jt}$$

The random investment cost shock v_{jt} helps to rationalize different investment decisions by observably similar firms. To our knowledge, this is one of the first empirical analyses that incorporates this feature in the context of continuous investment.

The optimal investment policy function $i_{jt}^*(s_{jt}, v_{jt})$ is implicitly defined by the first order condition in equation (3.6). Our goal is to flexibly estimate $i_{jt}^*(s_{jt}, v_{jt})$. Under reasonable assumptions, one can show that the optimal investment is monotonically decreasing in v_{jt} : firms with more favorable (smaller) cost shocks invest more, all else equal.²⁰ As a result, conditioning on s_{jt} , the j^{th} quantile of v_{jt} corresponds to the $(100 - j^{th})$ quantile of i_{jt} in the data. As shown in Bajari *et al.* (2007), we can recover the

²⁰One sufficient condition for monotonicity is that the value function has increasing differences in investment and the negative of the investment shock.

optimal investment policy function $i_{jt}^*(s_{jt}, v_{jt})$ as follows:

$$\begin{aligned} F(i|s_{jt}) &= Pr(i_{jt}^* \leq i|s_{jt}) = Pr(v_{jt} \geq i^{*-1}(s_{jt}, i)|s_{jt}) = Pr(v_{jt} \geq v|s_{jt}) \\ &= 1 - F_v(v|s_{jt}) \end{aligned}$$

$$\text{which implies } i_{jt}^*|s_{jt} = F^{-1}(1 - F_v(v|s_{jt})) \quad (3.11)$$

where $F(i|s_{jt})$ denotes the empirical distribution of investment given the state variables and F_v is the distribution of v . The data requirement for estimating this conditional distribution non-parametrically increases dramatically with the number of state variables. We make the simplifying assumption that the investment cost shock v_{jt} is independent of observed state variables s_{jt} and is additive:

$$i_{jt}^* = h_1(s_{jt}) + h_2(v_{jt})$$

where both $h_1(s_{jt})$ and $h_2(v_{jt})$ are unknown functions to be estimated. Moreover, since the distribution of v_{jt} cannot be separately identified from $h_2(v_{jt})$ non-parametrically, we assume that v_{jt} is distributed standard normal. We first flexibly regress observed investment on the state variables to obtain an estimate of $h_1(s_{jt})$. Then we treat $i_{jt}^* - \hat{h}_1(s_{jt})$ as the relevant data and estimate function $h_2(\cdot)$ using equation (3.11).

We do not incorporate divestment in our analysis. Compared to the massive investment undertaken by Chinese shipyards, divestment is much less common and an order of magnitude smaller.²¹ Modeling the level of divestment with irreversible adjustment costs (firms only recoup a fraction of the nominal value of their capital goods when they sell them) introduces a kink in the cost function and makes the value function non-differentiable, which raises considerable computational challenges. As a result, we abstract away from formulating the magnitude of divestment.

²¹The aggregate divestment is about 12% of the aggregate positive investment in the industry. We also drop 5% outliers with investment exceeding RMB 250 million or capital stocks exceeding RMB 4 billion.

We perform two robustness checks to address the issue that investment is non-negative. The first is a Tobit model that assumes $h_2(v_{jt})$ is normally distributed. The second approach estimates $h_2(v_{jt})$ non-parametrically based on the Censored Least Absolute Deviation estimator (CLAD) proposed by Powell (1984). Appendix C.1.4 provides more details.

State Transition Process Some of the state variables included in s_{jt} , such as the province and ownership status, are fixed over time. The transition process for age is deterministic. Capital (k_{jt}) depreciates at a common rate δ , so that

$$k_{jt+1} = (1 - \delta)k_{jt} + i_{jt}$$

We calibrate δ to 2.3% quarterly (Brandt *et al.*, 2012), reflecting China's high interest rates over our sample period. Similar to capital, the firm's backlog in period $t + 1$ is determined by orders and deliveries in period t . We assume backlog at time $t + 1$ satisfies an AR(1) process: $b_{jmt+1} = (1 - \delta_{bm})b_{jmt} + q_{jmt}$, and calibrate δ_{bm} based on average deliveries.²²

Finally, we need to specify firm beliefs over the transition process of steel and ship prices. The steel price, which is perceived as exogenous to the industry, follows an AR(1) process. The equilibrium price for each ship type is a complicated object, determined by the intersection of aggregate demand and supply. Following other work in the literature (e.g. Aguirregabiria and Nevo 2013), we model shipyards' beliefs about ship prices as an AR(1) process. This is a behavioral assumption: firms do not follow the production decisions of hundreds of rivals to predict future ship prices. They instead use the AR(1) process as a heuristic rule. The introduction of the Chinese

²²The quarterly depreciation rate for backlog, δ_{bm} , equals 6.8% for bulk, 6.3% for tankers, and 6.2% for containers.

government policies presents a permanent and unanticipated shock to the industry, which can potentially change the evolution of firm beliefs over prices. To capture this, we allow the AR(1) process to differ before and after 2006 when the policies came into effect.

Value Function Approximation Armed with estimates of the policy functions and state transitions, we now turn to the value function. We assume that the scrap value ϕ_{jt} is distributed exponentially with parameter $1/\sigma_\phi$ and obtain the ex ante value function (i.e. prior to the realization of ϕ_{jt}) as follows:

$$\begin{aligned} V(s_{jt}) &\equiv E_\phi V(s_{jt}, \phi_{jt}) = E_\phi [\pi(s_{jt}) + \max\{\phi_{jt}, CV(s_{jt})\}] \\ &= \pi(s_{jt}) + p^x(s_{jt}) E(\phi_{jt} | \phi_{jt} > CV(s_{jt})) + (1 - p^x(s_{jt})) CV(s_{jt}) \\ &= \pi(s_{jt}) + p^x(s_{jt}) \sigma_\phi + CV(s_{jt}) \end{aligned} \quad (3.12)$$

where we use the fact that $E(\phi | \phi > CV) = \sigma_\phi + CV$, as shown in Pakes *et al.* (2007). $\pi_{jt}(s_{jt})$ and $p^x(s_{jt})$ denote firms' static profit and exit probability, respectively, and $CV(s_{jt})$ denotes the firm's continuation value as defined in equation (3.4).

The ex-ante value function in our context is smooth and can be approximated arbitrarily well by a set of B-spline basis functions of the state variables:

$$V(s_{jt}) = \sum_{l=1}^L \gamma_l u_l(s_{jt})$$

where $\{u_l(s_{jt})\}_{l=1}^L$ are basis functions and $\{\gamma_l\}_{l=1}^L$ are coefficients to be estimated. This approach has several advantages. First, it avoids discretization and approximation errors therein when the state space is large. Second, replacing an unknown function with a finite set of unknown parameters substantially reduces the computational burden.²³ Third, the accuracy of the value function approximation can be controlled via

²³ Another popular approach to calculate the value function is via forward simulation. The computational

appropriate choices of the basis functions and is directly benchmarked by the violation of the Bellman equation (3.12).²⁴

In order to further simplify the computational burden, we utilize the fact that a number of states enter the firm's marginal cost linearly. As such, instead of keeping track of each state separately, we collapse them into a one-dimensional index, $\bar{s}_{jt}$, using the estimated production cost coefficients. This index measures firms' observed cost efficiency: a higher $\bar{s}_{jt}$ is associated with a lower marginal cost and a higher variable profit. Appendix C.1.6 provides more details on the value function approximation and describes how we construct the state space and estimate the approximating parameters γ .

Second Stage

Investment and Exit We estimate the dynamic parameters $\theta^i \equiv \{\sigma_\phi, c_1, c_2, c_3, c_4\}$ via MLE. Our sample likelihood includes both the likelihood for exit decisions and the likelihood for investment decisions. The log-likelihood for exit is:

$$\sum_{j,t} \log(f(\chi_{jt})) = \sum_{j,t} \left[(1 - \chi_{jt}) \log(1 - e^{-\frac{CV(s_{jt}; \gamma)}{\sigma_\phi}}) - \chi_{jt} \frac{CV(s_{jt}; \gamma)}{\sigma_\phi} \right]$$

where $\chi_{jt} = 1$ if firm j exits in period t and $f(\chi_{jt})$ is the associated likelihood.

Optimal investment $i_{jt}^* = i^*(s_{jt}, v_{jt})$ is defined by the first order condition in (3.6). By construction, when $i^*(s_{jt}, v_{jt})$ is positive, it is strictly monotonic in v_{jt} . Assuming it

burden of our approach is comparable to forward simulation when the policy function is linear in parameters.

²⁴We refer the interested reader to supplemental material in Barwick and Pathak (2015) and Kalouptsi (2018) for Monte Carlo evidence on the performance of value function approximations.

is also differentiable, the likelihood of investment can be written as follows:²⁵

$$g(i_{jt}) = \begin{cases} \frac{f_v(v_{jt})}{|i'(v_{jt})|} & \text{if } i^*(s_{jt}, v_{jt}) > 0 \\ Pr\left(\left[\beta \frac{\partial E(V(s_{jt+1}; \gamma)|s_{jt}, i)}{\partial i} - \frac{\partial C^i(i, v_{jt})}{\partial i}\right]_{i=0} \leq 0\right) & \text{if } i^*(s_{jt}, v_{jt}) \leq 0 \end{cases}$$

where in the first row, $f_v(v_{jt})$ is the density of the cost shock v_{jt} and $|i'(v_{jt})|$ is the absolute value of the derivative of $i^*(s_{jt}, v_{jt})$ with respect to v_{jt} .

Since the scrap values ϕ_{jt} and investment shocks v_{jt} are assumed independent, the joint log-likelihood for exit and investment decisions is the sum of the two respective log-likelihoods. We maximize the sample log likelihood subject to the constraint that the Bellman equation (3.12) is satisfied (see Appendix C.1.6 for details):

$$\begin{aligned} \max_{\theta^i} L &= \sum_{j,t} \log(f(\chi_{jt}; \theta^i)) + \sum_{j,t} \log(g(i_{jt}; \theta^i)) \\ \text{s.t. } V(s_{jt}; \theta^i) &= \pi(s_{jt}) + p^x(s_{jt})\sigma_\phi + CV(s_{jt}; \theta^i) \end{aligned} \quad (3.13)$$

Entry Cost Parameters Estimating the distribution of entry costs is straightforward once the investment cost and scrap value parameters are known. A potential entrant enters if their value of entry exceeds the random draw of the entry cost:

$$\kappa_{jt}(T_t) \leq VE(s_{jt}) \equiv E[-C^i(k_{jt+1}) + \beta E[V(s_{jt+1})|s_t]]$$

We first construct the value of entry VE_{jt} using the dynamic parameter estimates and then estimate the mean entry costs using the observed entry decisions, via MLE.

One issue with the entry cost estimation is how to treat firms' initial capital stock. As the initial capital is an order of magnitude larger than observed post entry investment, our investment model cannot rationalize it as an ordinary investment

²⁵The necessary condition for differentiability is that the value function is twice differentiable in investment, which holds since the value function is approximated using smooth spline basis functions.

decision.²⁶ Therefore, we assume that the cost of the initial capital equals $C(k_{t+1}) = c_1 k_{t+1} + c_4 T_t k_{t+1}$, where k_{t+1} is drawn from the observed distribution of initial capital stocks. Essentially, we assume that firms face no adjustment costs when choosing their initial capital so that $c_2 k_{t+1}^2$ and $c_3 v_{t+1} k_{t+1}$ do not apply.

3.5 Results

This section follows closely the sequence in Section 3.4. Section 3.5.1 presents estimates of the static parameters (demand and production costs). Section 3.5.2 discusses estimates of the policy functions and state transition process. Section 3.5.3 reports the dynamic parameter estimates (investment cost, scrap values, as well as entry costs).

3.5.1 Static Demand and Production Cost Estimates

Demand Appendix Table C.2 reports estimates of the demand curve (3.9). Column (1) presents the simplest specification where the only demand shifter is the type-specific freight rate. Column (2) adds type-specific demand shifters. Column (3) further controls for a time trend, while Column (4) allows the time trend to differ before and after 2006. In all specifications, we allow for a different price coefficient before and after 2006, to capture changes in the slope of the demand curve after the introduction of Chinese subsidies. Given the limited number of observations for each type-specific aggregate demand, we restrict the price coefficient post 2006 and the coefficient on backlog to be

²⁶The mean capital stock upon entry is RMB 125 million, compared to the average investment of RMB 18.5 million.

the same across types. We instrument prices using the steel price and steel production and we estimate the demand system using GMM. Adding demand shifters improves the fit, though time trends appear to matter little.²⁷ As such we use Column (2) as our preferred specification.²⁸

Demand becomes less elastic post 2006. According to our preferred specification (Column 2), the price elasticity prior to 2006 was 1.8 for bulk and tanker, and 3.4 for containers. It fell to 0.3 for bulk, 0.6 for tanker, and 1.7 for containers post 2006.²⁹ As expected, demand is also responsive to backlog (which affects the future competition that shipowners face): a 1% increase in the backlog leads to a 1% decrease in the quantity of new ships demanded. The remaining shifters have the expected sign.

Production Cost Table 3.2 shows the estimated marginal cost parameters for Chinese yards for each ship type (standard errors are computed from 500 block bootstrap samples). We allow the key parameters that characterize the curvature of the production cost to be type specific (the coefficients on quantity, capital, backlog, and steel price), but restrict the coefficients on subsidy dummies and shipyard attributes to be the same across ship types, due to the large number of parameters. There are intuitive reasons for these restrictions. For example, the effect of scale economies from holding a large backlog, the return to capital (which proxies for capacity), and input intensity are likely

²⁷In addition, the time trends pose a challenge to the stationarity assumption of our dynamic setup and are difficult to deal with when extrapolating beyond our sample period in the counterfactual analyses.

²⁸Our results are similar if we use a different specification of the demand curve or vary the demand elasticity.

²⁹Demand elasticity for new ships, which are durable goods, is driven by complicated dynamic considerations that include the composition of the existing fleet, the expected number of new ships to be delivered in the near future, as well as the beliefs about future freight rates and fuel costs. Hence, we have no prior as to whether it should increase or decrease post 2006.

to be different across ship types. On the other hand, it is not unreasonable to assume that subsidies are not earmarked for a particular ship type, as firms can produce different kinds of ships depending on the prevailing market conditions.

As the Chinese policies came into effect in 2006 and underwent major changes in 2009, we include separate dummies for Chinese yards in 2006-2008 and from 2009 onwards. The production subsidy between 2006-2008 is estimated to be 1,510 RMB/CGT, which is 10-13% of the average price. The subsidy from 2009 onwards is slightly smaller, at 1,380 RMB/CGT.³⁰ Though our estimation method, sample period, and industry coverage are different from those in Kalouptsidi (2018), the estimated production subsidy is of a similar magnitude (with ours being slightly smaller), which is reassuring.

The parameter β_q captures the increase in marginal cost (in 1000 RMB/CGT) from taking an additional order of 100,000 CGT. The larger β_q is, the more convex the cost function is, and the less responsive supply is to changes in prices. On average, a 10% price increase causes bulk production to increase by 21%, tanker production by 28%, and container production by 22%. Higher capital is associated with a lower marginal cost of production, though at a diminishing rate (the coefficient on capital squared is positive). Increasing capital by RMB 100 million for an average firm with a capital of RMB 400 million reduces the annual marginal cost of production by 2.1% for bulk, 1.8% for tanker, and 1.7% for container. To put these numbers into context, the average firm's per-period profits would decline by 19% if its capital stock were halved.

Moreover, we find strong evidence of economies of scale in production with respect

³⁰In robustness checks, we estimate the production subsidies separately for each region. They are higher in Jiangsu and Liaoning than in Zhejiang and the rest of China, although the differences are statistically insignificant.

Table 3.2: Cost function estimates

	Bulk		Tanker		Container	
Type-specific	Coefficient	T-stat	Coefficient	T-stat	Coefficient	T-stat
β_q	7.34	9.52	13.60	5.54	9.69	5.63
σ_ω	8.49	10.43	14.40	7.08	12.14	5.71
Constant (1000 RMB/CGT)	19.26	15.88	36.58	9.18	32.30	8.39
Steel Price (1000 RMB/Ton)	1.55	7.49	1.10	3.04	0.63	1.65
Capital (bill RMB)	-2.43	-2.96	-2.61	-1.80	-2.19	-2.01
Capital ²	0.19	0.83	0.06	0.25	0.06	0.32
Backlog	-1.56	-5.29	-4.44	-5.04	-2.88	-3.34
Backlog ²	0.07	4.04	0.24	3.43	0.18	1.97
Backlog of Other Types	0.13	0.94	0.35	1.65	0.46	2.66
Common						
2006-2008	-1.51	-2.62				
2009+	-1.38	-2.37				
Large firms	-3.85	-6.97				
Jiangsu	-2.64	-4.75				
Zhejiang	-1.42	-2.80				
Liaoning	-1.87	-2.05				
CSSC/CSIC	-0.77	-1.20				
Private	0.14	0.30				
Foreign JV	-0.78	-1.45				
Age	0.18	3.14				
N	4886		4977		2504	

Note: Standard errors bootstrapped using 500 bootstrap samples.

to the backlog: it is cheaper to produce multiple ships at the same time. The effect of the backlog on marginal cost is sizable: increasing the backlog by 100,000 CGT reduces the marginal cost of production by 11% to 27% on average across ship types. As backlogs continue to increase, capacity constraints become binding and drive up marginal costs, as reflected in the positive coefficient (though much smaller in magnitude) on backlog squared.

Firms located in Jiangsu, Liaoning, and Zhejiang provinces (the major shipbuilding regions in China) have lower marginal costs, by 18-22% for Jiangsu, 13-16%

for Liaoning, and 10-12% for Zhejiang. The (additional) effect of ownership is limited. While private firms have the highest marginal costs, followed by small SOEs, CSSC/CSIC owned SOEs, and finally foreign JVs, none of these coefficients is statistically significant. As shipyards age, their marginal cost increases by about 1% each year. Finally, increases in the steel price raise the marginal cost for all types, as expected.

The fixed cost calibrated from the average of the NBS accounting data equals RMB 15 million per quarter, accounting for 15% of the industry profit on average. Thus setting it to zero, as is commonly done in the literature, would significantly overestimate per-period profits accruing to firms.

Robustness Our baseline specification (Table 3.2) assume price-taking behavior on the part of shipyards. Appendix Table C.8 illustrates how our estimates of production cost change when we relax this assumption and instead allow firms to compete in quantities. The resulting cost function estimates are similar to our baseline estimates, which is unsurprising given the unconcentrated nature of the industry.

We estimate costs separately for each country in the baseline specification, mostly because we only observe the capital stock for Chinese shipyards. Appendix Table C.9 displays parameter estimates when shipyards from all three countries are pooled together; this specification amounts to a differences-in-differences estimator. We set the capital variable to zero for Japanese and South Korean yards and use country dummies to control for these capital costs. The key coefficients are qualitatively similar, though the subsidy for Chinese yards during 2006-08 is estimated to be somewhat larger relative to the baseline specification. We prefer the baseline specification, which allows more flexibility in capturing production differences across countries and delivers a more

conservative estimate of the subsidies. Since the baseline estimates deliver smaller subsidies, our counterfactual results in Section 3.6 can be viewed as lower bounds.³¹

Finally, we look at whether there is evidence of learning-by-doing among Chinese shipyards (Benkard, 2004). We examine both within-firm and industry-wide learning-by-doing by allowing the marginal cost of production to depend on a firm's past production as well as the industry cumulative past production. The results are illustrated in Appendix Table C.10. Despite the potential upward-bias of over-estimating the spillover effects in the absence of suitable instruments, our estimates suggest that there is no evidence of learning-by-doing: marginal costs tend to *increase* rather than decrease in past production. This is consistent with industry reports that the technology for producing ships, especially bulk and tankers, has been around for a while and is mature. We have also estimated production costs while allowing for a time trend, and on a subsample that drops new shipyards that entered after the policy announcement, and the results are broadly similar (see Appendix Table C.11). We hence take the baseline estimates from Table 3.2 to the dynamic estimation and counterfactual analyses.

3.5.2 Dynamic Parameters: First Stage

Investment policy function Appendix Table C.14 reports the estimates for the investment policy function using OLS, Tobit with $h_2(v)$ normally distributed, and CLAD that does not impose a distributional assumption on v . Our preferred specification is the OLS regression, which delivers the highest model fit. Investment increases in ship prices and decreases in the steel price. Firms with higher $\bar{s}_{jt}$ (i.e., more productive) invest more all else equal. As expected, the coefficients for both the 2006-

³¹Using the cost estimates that pool data from all three countries leads to qualitatively similar counterfactual results.

08 and 2009+ policy dummies are positive. Moreover, investment is hump-shaped with respect to capital: it initially increases in the capital stock, reaches a peak when capital equals RMB 1-1.5 billion, and then falls.

Exit policy function We estimate the exit policy function via a probit regression. Appendix Table C.13 presents two sets of estimates using linear terms of all states as well as capital squared, with and without region fixed effects. Firms with higher $\bar{s}_{jt}$ are less likely to exit, which is intuitive as $\bar{s}_{jt}$ is a measure of firm profitability. Exit probabilities are lower when subsidies are in place.

3.5.3 Dynamic Parameters: Second Stage

Investment and exit Table 3.3 reports the investment cost estimates. Following the literature on investment (Cooper and Haltiwanger, 2006), we assume that the unit investment cost is equal to one ($c_1 = 1$).³² Between 2006-2008, the subsidy was 0.25 RMB per RMB of investment, implying that 25% of the per-unit cost of investment (excluding adjustment costs) is subsidized. Post 2009, the subsidy nearly doubles and jumps to 0.49 RMB per RMB of investment, which helps rationalize the elevated investment post the financial crisis with plummeting ship prices. In addition, the increase in subsidies post 2009 is consistent with the government policy change that shifted the focus towards consolidating the industry and supporting existing firms.

The coefficient on quadratic investment, c_2 , is both economically and statistically significant. On average, adjustment costs account for 28% of total investment costs and exceed 50% for large investments over RMB 50 million. In addition, an estimated

³²Monte Carlo evidence indicates that it is difficult to identify all the cost parameters in equation (3.7).

Table 3.3: Estimates of investment cost and scrap value parameters

	Coeff.	T-stat
σ_ϕ	0.69	11.76
c1	1.00	
c2	21.72	10.57
c3	1.55	8.27
$c4_{2006-08}$	-0.25	-1.89
$c4_{2009+}$	-0.49	-4.07
N	4286	

Note: Standard errors bootstrapped using 500 block bootstrap samples.

large value of c_3 implies that the shock v_{jt} plays an important role in explaining investment. Finally, the mean of the scrap value distribution is estimated to equal RMB 0.69 billion. This is significantly lower than the estimated value of a firm, $V(s_{jt})$, which is around three to four billion RMB, as exit is a rare event and occurs in only 1% of the observations.

Appendix Figure C.2 plots the distribution of the observed and the simulated investment. These two distributions are reasonably similar, though the actual distribution has a heavy-tail of large investments and fewer medium-sized investments. Finally, Appendix Table C.3 compares the actual number of exits with the model's estimates. Our model predicts fewer exits post 2006, with a total of 39 exits compared to a total of 48 exits observed in the sample.

Entry cost estimates Table 3.4 reports estimates for κ_{jt} , the mean entry costs, for period before 2006, between 2006 and 2008, and post 2009 respectively. Our estimated number of entrants is reasonably close to the actual number of entrants in each policy period (Appendix Table C.4). Given the different number of entrants across provinces (Zhejiang has the highest number of entrants at 95), we estimate the entry cost separately for Liaoning, Jiangsu, Zhejiang, which collectively account for 70% of new shipyards,

and the rest of China.³³ The entry cost estimates range from RMB 25 billion to 91 billion prior to 2006. Conditional on entering, the average entry cost paid is RMB 2.3 billion, close to a shipyard's accounting value.³⁴ Given the unprecedented entry boom from 2006 to 2008, it is not surprising that we find substantial entry subsidies, with the fraction of entry costs that is subsidized varying from 49% in Liaoning to 64% in Jiangsu. Entry costs increased substantially in 2009 when the entry moratorium was put in place.

Table 3.4: Entry Cost Distribution (Mean), billion RMB

	κ_{pre}	$\kappa_{post,06}$	% of pre costs	$\kappa_{post,09+}$	% of pre costs
Jiangsu	60	22	36%	69	114%
Zhejiang	91	37	41%	194	214%
Liaoning	56	29	51%	-	-
Other	25	10	38%	44	172%

Note: κ_{pre} : mean of the entry cost distribution prior to 2004 for Zhejiang, and prior to 2006 for Jiangsu, Liaoning and Other regions. $\kappa_{post,06}$: mean of the entry cost distribution between 2004 and 2008 for Zhejiang, between 2006 and 2008 for Jiangsu, Liaoning and Other regions. $\kappa_{post,09+}$: mean of the entry cost distribution from 2009 onwards. The number of potential entrants, $\bar{N}$, is assumed to equal twice the maximum number of potential entrants ever observed in a region.

A common challenge in estimating entry costs is setting the number of potential entrants, as this is rarely observed. Following the literature (Seim, 2006), we assume that the number of potential entrants in a region in any quarter is larger than the maximum number of entrants ever observed in that region, specifically, twice the maximum number of observed entrants. At this level, the entry rate is only 6.8% and thus leads to high estimated entry costs. We have estimated the entry cost distribution under

³³Entry subsidies are assumed to begin in 2004 for Zhejiang, when it identified shipbuilding as a pillar industry, and in 2006 for all other provinces. This is consistent with the empirical pattern that entry peaked earlier in Zhejiang than the rest of the country. The ship type produced by an entrant is a random draw from the observed distribution of product lines for entrants in that province.

³⁴See <http://www.jiemian.com/article/1483665.html> and http://www.wuhu.com.cn/compay_mod_file/news_detail.php?cart=3&id=595 for news articles that report the book value of shipyards.

alternative assumptions (e.g. the maximum number of entrants ever observed, or a large number such as 20 and 40). The higher the number of potential entrants, the higher the estimated κ_{jt} . While the estimates for κ_{jt} vary, the estimated entry cost paid upon entering, as well as the entry subsidies are remarkably robust, as they are essentially determined by the actual number of entrants observed in the sample.

3.6 Evaluation of China's Industrial Policies

In this section, we evaluate the long-term implications of China's industrial policy. Doing so necessitates simulating the world shipbuilding industry for a long period of time, as both entry and investment have dynamic consequences – the accumulated capital remains productive and new firms can continue operation long after the policy ends.³⁵ All monetary values reported in this paper are discounted and deflated to the 2006 RMB.

Our simulation begins in 2006, when the Chinese government started subsidizing its domestic industry, and ends in 2050, a period long enough to evaluate a policy's dynamic impacts, though results are similar if the simulation ends in 2099 (or between 2050 and 2099). In each counterfactual scenario, we turn on and off the subsidies as needed.³⁶ In each period, Chinese firms make optimal production, investment, exit and entry decisions, taking both prices and government policies as given. Japanese and

³⁵Production subsidies also have dynamic consequences through backlogs that affect future costs of production, though these effects disappear within a few years when backlogs are converted to deliveries.

³⁶For each counterfactual analysis, we conduct 50 simulations. Results are nearly identical with 100 simulations. The discount rate is 0.02 per quarter, or 0.08 annually, reflecting the high interest rates in China (averaging 6% from 1996 to 2018). Using different discount rates leads to qualitatively similar results.

South Korean firms make production decisions, but do not invest or enter/exit, since there is limited capacity expansion or entry/exit among these firms as discussed above; we also lack the relevant firm-level data. Equilibrium prices are determined each period by the intersection of the industry demand and supply curves. Appendix C.2 contains more details on the implementation.

In Section 3.6.1, we first quantify how the policy affected the actual evolution of the domestic and global industry between 2006 and 2013. Then, we turn to the main exercise and assess the long-run performance of different policy instruments. We also discuss various aspects of policy design, such as the timing of subsidies and targeting. Section 3.6.2 evaluates the consolidation policy. Finally, in Section 3.6.3, we examine potential rationales for this industrial policy intervention.

3.6.1 Evaluation of Industrial Policy

The Impact of Industrial Policy on Industry Evolution Perhaps not surprisingly, subsidies had a significant impact on the realized evolution of every outcome of interest: China's market share, total ship production, ship prices, entry and exit, investment, profits, concentration and capital utilization.

Total (discounted) subsidies handed out to Chinese shipbuilders between 2006 and 2013 were close to RMB 460 billion (or \$77 billion), with the lion's share going to entry subsidies (RMB 297 billion), followed by production subsidies (RMB 125 billion) and investment subsidies (RMB 37 billion).³⁷ The subsidies are massive in comparison to

³⁷While entry subsidies are large in magnitude, they are consistent with a back-of-the-envelope calculation: entry subsidies induced the entry of 80-90 additional firms and each firm is worth a few billion RMB.

the size of the domestic industry, whose revenue was around RMB 1360 billion during the same period.

These subsidies increased China's world market share during 2006-13 by nearly 40%. The ascent in market share is most pronounced for bulk ships, since a large fraction of the new shipbuilders produce bulkers and the cost advantage enjoyed by Japanese and South Korean firms is narrower for bulkers. In the absence of subsidies, China's production would still increase in absolute value in response to the higher demand during the boom, though the increase would have been much less pronounced.

In absolute terms, only one-fourth of China's increased production translated into higher world industry output. The remaining three-fourths constitute business-stealing, whereby Chinese production expanded at the expense of competing firms in other countries. As a consequence of Chinese subsidies, South Korea's market share decreased from 47% to 38% and Japan's market share from 24% to 21% during the 2006-2013 period, with profits earned by shipyards in these two countries falling by RMB 140 billion.

The rising global supply induced by the subsidies led to substantial reductions in global ship prices: the price of bulk ships, oil tankers, and containers fell by 8.2%, 6.2%, and 3.1% from 2006 to 2008, respectively (Table C.5). The price effect is most significant for bulkers, as Chinese shipyards account for a bigger market share in bulk and demand for bulkers is the most inelastic. As the effect of past subsidies accumulates through an increased production capacity for existing firms and a larger number of new firms, the price drop became more pronounced from 2009 to 2013 and reached 16.5% for bulk, 10.6% for tankers, and 3.7% for containers. Because of the reduction in ship prices, world shipowners benefit by RMB 230 billion from Chinese subsidies, though only a small proportion of these gains accrues to Chinese shipowners as they account

for a small fraction of the world fleet.³⁸

Figure 3.5 compares the number of Chinese firms by year with and without subsidies. Government support more than doubled the entry rate: 148 firms enter with subsidies vs. 65 without subsidies from 2006 to 2013. It also depressed exit (37 firms exit vs. 46) and changed the composition of exiting firms. With subsidies, a bigger fraction of exiting firms come from those that entered during the policy period, partly because subsidies attract small and relatively inefficient firms.


Figure 3.6 illustrates the striking effect of subsidies on investment, which skyrocketed post 2006. Total investment during 2006-2013 is RMB 85 billion with subsidies, compared to RMB 32 billion without subsidies.

The policy led to increased fragmentation. Entry subsidies induce entry of small inefficient firms. In addition, production and investment subsidies boost firms' variable profit and retain unprofitable firms that should have exited. As a result, the industry became much more fragmented post the policy period. China's domestic Herfindahl-Hirschman index (HHI) plummeted from 1,200 in 2004 to less than 500 in 2013 with a significantly lower 4-firm concentration ratio (Appendix Figure C.3). Despite a sizable increase in China's overall production, capacity utilization was much lower, particularly when demand was low post 2009. If China had not subsidized the shipbuilding industry, the ratio of production to capital (which proxies for capacity utilization) would have been 20% higher during the 2009-2013 recession.

Finally, total subsidies handed out between 2006-13 led to an estimated RMB 45 billion increase in the discounted lifetime net profits for Chinese shipbuilders. In other words, compared to the size of the intervention, profit gains are modest at best.


³⁸ According to Clarkson World Shipyard Monitor, orders by Chinese shipowners have been growing but still account for under 10% of world orders in 2010-2013.

Figure 3.5: No of firms, with and without subsidies



Note: Total number of firms in the case of all subsidies (as in the data) and a counterfactual case with no subsidies.

Figure 3.6: Investment, with and without subsidies



Note: Total investment in the case of all subsidies (as in the data) and a counterfactual case of no subsidies.

Long-run Performance of Policy Instruments Subsidies led to a substantial increase in the number of firms, capacity, and production, and also resulted in lower prices, lower capital utilization and a more fragmented industry structure during our sample period. We now turn to the assessment of the long-run performance of different policy instruments, in particular the relative efficacy of production, investment, and entry subsidies, and search for general lessons that can be applied in other contexts. Since China's domestic consumer surplus accounts for a small fraction of the world consumer surplus as discussed earlier, and is modest compared to the industry profit, we only briefly discuss the effect on consumer surplus when it is relevant and focus instead on industry outcomes, such as output and profits.

In order to compare different types of policies, we carry out five counterfactual exercises with different subsidies in place: all subsidies (as in the data), only production subsidies, only investment subsidies, only entry subsidies, and no subsidies. When simulating the industry beyond 2013, we assume that the 2013 policy environment is propagated to the end of our simulation period unless otherwise noted. For example, in the scenario with all subsidies, entry subsidies run from 2006 to 2008, whereas

production and investment subsidies run from 2006 to the end.³⁹

The results are summarized in Table 3.5, which reports the discounted sums of industry revenue and profit in China, as well as the magnitude of different subsidies. The last two rows, “ Δ Revenue/Subsidy” and “ Δ Net Profit/Subsidy”, constitute different measures of the effectiveness of the subsidies. “ Δ Revenue” is the revenue difference between subsidies and no subsidies. We report the ratio between the revenue increase and the subsidy cost to evaluate the effectiveness of subsidies at promoting industry revenue. This is of interest, as China’s official government documents explicitly state production targets for the domestic shipbuilding industry. “ Δ Net Profit” is the difference in net profit which equals revenue plus the scrap value upon exiting, minus the costs of production, investment, and entry. We use “ Δ Net Profit/Subsidy” to measure the *gross rate of return* of different subsidies. A rate lower than 100% indicates that the cost of subsidies exceeds the net benefits to the domestic industry.⁴⁰

When all subsidies are in place, this policy mix is highly ineffective, as reflected by the rate of return being merely 16%. When each policy is in place in isolation, the return is 53% for production subsidies, 78% for investment subsidies, and 24% for entry subsidies, respectively. This implies that the distortions induced by multiple subsidies are convex: i.e. the combination of all policies yields a considerably lower return compared to each policy in isolation. For example, entry subsidies lower the entry threshold and thus attract inefficient entrants. With the introduction of production and investment subsidies, the number of firms in operation is further inflated due to

³⁹Another option is to shut down production and investment subsidies post 2013. This would constitute an unanticipated policy shock to firms and is shown in our simulations to produce lower returns.

⁴⁰The discussion below does not take into consideration the cost to finance these subsidies. Estimates from Ballard *et al.* (1985) suggest that collecting one dollar of government revenue costs 17 to 56 cents in the U.S. Including the cost to finance subsidies will further drive down the rates of return.

Table 3.5: Comparison of Different Subsidies

	All Subsidies	Only Production	Only Investment	Only Entry	No Subsidies
Lifetime Revenue 2006-	2273	2048	1760	1801	1669
Lifetime Profit 2006-	831	768	602	578	558
Production Subsidy	249	210	0	0	0
Investment Subsidy	83	0	43	0	0
Entry Subsidy	302	0	0	171	0
Δ Revenue/Subsidy	95%	180%	214%	78%	
Δ Net Profit/Subsidy	16%	53%	78%	24%	

Note: Each element in the table refers to a discounted sum from 2006 to 2050, averaged across all simulations. For example, “Lifetime Revenue (Profit) 2006-” refers to the discounted sum of revenue (net profit) earned by firms from 2006 to 2050. “ Δ Revenue/Subsidy”: the discounted sum of revenue in the scenario minus the discounted sum of revenue with no subsidies, divided by the discounted sum of subsidies. “Net Profit”: Revenue-Production Cost-Investment Cost+Scrap Value-Entry Cost. “ Δ Net Profit/Subsidy” equals the discounted sum of net profits in the scenario minus the discounted sum of net profits with no subsidies, divided by the discounted sum of subsidies. In Column “Only Production”, we maintain the same production subsidy as in the baseline estimation, but shut down entry and investment subsidies. Columns “Only Investment” and “Only Entry” are similar. Government policy from 2014 onward remains the same as in 2013.

subsidized revenue. This drives down the rate of return and makes the subsidies more distortionary in per-dollar terms.

An important factor contributing to the low returns are fixed costs. Indeed, firms incur fixed costs to stay in business even when they receive no orders from buyers. In volatile industries with cycles of booms and busts, this tends to be a common occurrence: as exit is irreversible, firms are willing to suffer temporary losses in expectation of higher demand in the future. If fixed costs were zero, the rate of return on subsidies would increase from 16% to 27%.

We now turn to the performance of each type of subsidy in isolation. If industry revenue is the object of interest, both production and investment subsidies are effective. On average, one RMB increase in the production and investment subsidy raises the industry’s revenue by RMB 1.8 and RMB 2.1, respectively. This might justify the

popularity of these subsidies in China, since quantity and revenue targets are often linked to local officials' promotion (Jin *et al.*, 2005). In addition, as discussed earlier, the government had set explicit output targets for ship production.

The investment subsidies appear less distortionary than the production subsidies (78% vs. 53%). However, this comparison is confounded by the larger magnitude of the production subsidies, as in our setting bigger subsidies are associated with more distortion. We reduce the per-unit production subsidies by 75% to make the total amount of these two subsidies comparable (RMB 43 billion). The return on these two subsidies is then nearly identical (see Appendix Table C.6). Production subsidies are slightly more effective at increasing revenue: the increase in revenue per RMB of subsidy is 229% for production subsidies, versus 214% for investment subsidies. In a similar vein, Aldy *et al.* (2018) find that wind farms claiming output subsidies produced 10-11% more power than wind farms claiming investment subsidies.

Entry subsidy is the least efficient policy instrument among the three by a large margin.⁴¹ It predominantly attracts small and high-cost firms that would not find it profitable to enter and operate in the absence of subsidies. The large number of additional entrants contributes little to industry profits, while it exacerbates excess supply and reduces ship prices. In contrast, the take-up rate for production and investment subsidies is higher among firms that are more efficient, receive more orders (higher backlogs), and are more likely to invest.⁴² In addition, production and

⁴¹It might be easier to finance entry subsidies than production or investment subsidies, which matters for the implementation of different policies. Unfortunately we do not have data to evaluate the importance of this concern.

⁴²Appendix C.3 uses a simple static model to illustrate that the rate of return is higher when taken up by efficient firms. In our empirical analysis, all four decisions can be distorted: entry, production, investment, and exit. Compared to efficient firms, inefficient firms are more likely to be distorted in all four margins with subsidies. In addition, efficient firms enjoy considerable economies of scale as a

investment subsidies increase backlogs and capital stocks that lead to economies of scale and drive down both current and future production costs.

Table 3.6 illustrates these points by decomposing subsidies that are taken up by firms above or below the median efficiency (as measured by $\bar{s}_{jt}$), for column one of Table 3.5. Subsidies are significantly more effective when given to productive firms, in terms of increased both revenue and profits. The subsidies accruing to productive firms has a net rate of return equal to 27%, while those going to unproductive firms is almost entirely wasted.⁴³ Note that production and investment subsidies are reasonably well-targeted, with between 70-84% of subsidies going to productive firms. In comparison, less than 50% of entry subsidies are taken up by productive firms.

Table 3.6: Effect of subsidies on productive and unproductive firms

	Unproductive firms	Productive firms
Lifetime Revenue 2006-	404	1866
Lifetime Profit 2006-	70	757
Production Subsidy	40	209
Investment Subsidy	26	58
Entry Subsidy	157	146
Δ Revenue/Subsidy	71%	111%
Δ Net Profit/Subsidy	-5%	27%

Note: “Unproductive” firms: firms with initial $\bar{s}_{jt}$ below the median (at the time the policy change first occurred); “productive”: firms with initial $\bar{s}_{jt}$ above the median. All rows are defined as in Table 3.5.

Our welfare analysis so far has abstracted away from changes in consumer surplus enjoyed by Chinese shipowners. Assuming that these gains are equal to 10% of total surplus gains (which is an upper bound), the rate of return increases from 16% to 24%.

result of larger backlogs and capital stocks relative to inefficient firms, which further widens the wedge in efficiency.

⁴³The rate of return for subsidizing unproductive firms is negative due to a general equilibrium effect: increased production drives down ship prices. Since these firms have high production costs, reduction in ship prices offsets potential gains in quantity produced.

For the total benefit of the subsidies to exceed the cost, China's share of the global consumer surplus from new ships would need to be over 94%.

Business Cycles and Industrial Policy Like many other industries, cycles of booms and busts are a fundamental feature of shipbuilding. A rich macro and public finance literature explores the optimal fiscal policy over the business cycle and generally recommends counter-cyclical fiscal policies, in order to smooth out intertemporal consumption (Barro, 1979), reduce the efficiency costs of business cycle fluctuations (Gali *et al.*, 2007), and increase long-run investment by lowering volatility (Aghion *et al.*, 2014). It is less well-understood, however, how industrial policy should be optimally designed in the presence of boom-and-bust cycles.

To explore whether the effectiveness of subsidies varies over the business cycle, we carry out two counterfactual simulations. The first simulation only subsidizes production and investment of all firms during the 2006-08 boom, while the second simulation only subsidizes production and investment during the 2009-13 bust. All subsidies are discontinued afterwards. The subsidy rates are calibrated so that the government spends the same amount in the two scenarios. This design allows us to explore the long-run implications of pro-cyclical vs. counter-cyclical policies.

Strikingly, subsidizing firms during the boom leads to a net return of only 28%, whereas subsidizing firms during the downturn leads to a much higher return of 77%, as shown in Table 3.7. What explains this large difference?

There are two main contributing factors: convex production and investment costs, and firm composition. In booming periods, the industry is operating close to full capacity. Further expansion is costly and entails utilization of high-cost resources. Firms that are already producing and investing may choose to engage in more rapid expansion

Table 3.7: Pro-cyclical vs. counter-cyclical industrial policy


	Subsidize during boom	Subsidize during recession
Lifetime Revenue 2006-	1765	1767
Lifetime Profit 2006-	598	613
Production Subsidy	34	35
Investment Subsidy	16	16
Δ Revenue/Subsidy	221%	223%
Δ Net Profit/Subsidy	28%	77%

Note: In Column “Subsidize during recession”, the government subsidizes production and investment during the recession of 2009-13, but offers no subsidy before 2009 or after 2013. In Column “Subsidize during boom”, the government only subsidizes during the boom of 2006-08. The subsidy rates during the 2006-08 boom are adjusted downwards to match the amount handed out during the recession. No entry subsidy is offered in either scenario. All rows are defined as in Table 3.5.

than is optimal, incurring large adjustment costs. During a bust, on the other hand, the industry operates well below capacity and many production facilities remain idle. Subsidies mobilize underutilized facilities, resulting in smaller distortions. The second contributing factor is the changing firm composition over the business cycle. Subsidies during a boom attract a higher fraction of inefficient firms, which pushes down the rate of return. As an illustration, Figure 3.7 plots the average $\bar{s}_{jt}$ (a measure of profitability) over time for these two scenarios. Subsidies in the boom is associated with a much lower average profitability than subsidies in the bust, as expected.

Despite the benefits of a counter-cyclical policy, the actual policy mix is overwhelmingly pro-cyclical: total subsidies of RMB 391 billion were handed out between 2006 and 2008, vs. RMB 68 billion between 2009 and 2013. This echoes a more general finding in the literature showing that developing countries typically use pro-cyclical fiscal policies (Frankel *et al.*, 2014), due to budget constraints, political considerations, etc. (Tornell and Lane, 1999; Barseghyan *et al.*, 2013).

Figure 3.7: Average firm cost-efficiency $\bar{s}_{jt}$ with subsidies during the boom vs. subsidies during the bust



Note: Average firm cost-efficiency, as captured by $\bar{s}_{jt}$, when subsidies are distributed during a boom vs. during a bust. Variable $\bar{s}_{jt}$ is defined in Appendix C.1.6.

3.6.2 Evaluation of Consolidation Policies

One explicit goal of China's industrial policy after the financial crisis is to facilitate consolidation and create large successful firms that can compete against international conglomerates. A crucial policy for achieving this objective was the 2013 *Shipbuilding Industry Standard and Conditions*, whereby the government announced a list of selected firms that meet the industry standard, the so called "White List". In this section, we ask the following questions. First, does the consolidation policy improve the return of subsidies, and if so by how much? Second, did the government choose the optimal set of firms for the White List?

Gains from Targeting In our first counterfactual exercise, we rank firms based on their expected variable profits ($E[\pi_{jt}]$) in 2013, and select the top 56 firms with the

highest profitability to form the “Optimal White List”.⁴⁴ These firms receive production and investment subsidies, while other firms receive no subsidies post 2013. We compare this policy to the one that subsidizes all firms after 2013, as well as the scenario with no subsidies.

As shown in Appendix Table C.7, directing subsidies towards the best set of firms generates *considerable* gains. The net rate of return for targeted production and investment subsidies is 85%, whereas the return is 34% when all firms are subsidized. This pattern holds across both measures of policy effectiveness (revenue and net profit), due to several reasons. First, subsidizing all firms encourages sub-optimal entry, while the White List policy only subsidizes existing firms and does not distort entry.⁴⁵ Second, subsidizing existing firms leads to lower exit than is socially optimal, but the set of firms on the White List is less likely to be subject to sub-optimal exit decisions, in contrast to other firms that are more prone to distortions. Lastly, as argued above, targeting productive firms leads to less distortion.

White List While subsidies are less distortionary when targeted towards efficient firms, it is unclear a priori whether the government targeted the optimal set of firms. Information asymmetries and regulatory capture might bias the process in favor of interest groups or “sunset sectors” (Lane, 2017).

To avoid confounding effects from subsidies and focus on the “White List” only, we discontinue all subsidies post 2013 and examine post-2013 profit for firms on the actual White List and firms in our “optimal White List” as constructed above. Note that our selection criterion focuses on short-run profitability and we may miss firms with the

⁴⁴Four out of sixty firms on the official White List cannot be matched to the rest of our datasets described below, hence we focus on the remaining fifty-six firms in our counterfactual analysis.

⁴⁵The effect on entry is negligible.

highest long-run profits. Thus this is a weak test: if the government chose firms with the highest long-run profitability, then their selection should do at least as well as the set of firms we choose.

As shown in Appendix Figure C.4, industry profits are lower with the actual White List (the dashed blue line) than our “optimal White List” (the solid red line). The difference in the long-run industry profits and revenue (the discounted sum from 2014 to 2050) is 14% and 10%, respectively, suggesting that the government did not choose the optimal set of firms in its White List. Out of the 56 firms chosen by the government, only 31 firms appear in our White List based on the short-run profitability. There appears a bias in favor of SOEs: 65% of firms selected by the government are SOEs, while 55% of our selected firms are SOEs.

3.6.3 Rationales for Industrial Policy

Our evaluation of China’s industrial policy in promoting the shipbuilding industry so far is mixed. The subsidies boosted production and investment, but contributed to the development of a fragmented industry, reduced the capital utilization rates, and to a large extent, were dissipated through sub-optimal entry/exit decisions. In this section, we assess traditional arguments in favor of industrial policies and evaluate the extent to which existing policies are effective in achieving these objectives.

A standard justification for subsidies is the presence of positive externalities (such as industry-wide learning-by-doing), as each firm produces less than what is socially optimal. As we discussed in Section 3.5.1, there is no evidence of significant spillover effects in this industry. This is consistent with industry reports that much of the production by Chinese shipyards occurs in product sectors with mature technologies,

where the scope for learning is limited.⁴⁶

Our analysis focuses on the shipbuilding sector and does not account for benefits to other sectors. While spillovers to downstream sectors provide a rationale for subsidizing upstream industries (Liu, 2019), it is unlikely a justification for subsidies in the shipbuilding industry. Three-quarters of the output from this industry is used for final consumption (China's 2012 Input-Output Table). Moreover, most of these ships are exported, reducing the share of benefits from subsidies that is captured domestically.

There are potential spillovers to upstream sectors. Intermediate inputs from other sectors account for 63% of the value of ships produced. Steel in particular is an important input. One might argue that shipbuilding subsidies are partially designed to boost demand for steel, a strategic sector that has been subject to many policy interventions. However, steel used in shipbuilding accounts for less than 1.5% of total steel produced (China's 2012 Input-Output Table).

Another justification for industrial policy is labor market consequences: subsidies could have welfare benefits if they increase employment or offset distortions that lead to depressed employment. Even in the grand scheme of things, total employment in shipbuilding and related industries (ship repairs and marine equipments, etc.) accounts for less than 0.5% of national employment, suggesting that any potential labor market benefits would be modest.

In the presence of market power, there are in principle strategic trade benefits from subsidizing industries that compete with foreign firms (Dixit, 1984; Krugman, 1986;

⁴⁶There might be technological 'catching-up' and learning among Chinese shipyards for producing the latest generation ships (e.g. large containerhips or LNG's), where most of the patents and 'know-how' are possessed by Japanese and South Korean firms. Unfortunately, there are few orders of these ships and we cannot directly test this.

Eaton and Grossman, 1986; Brander, 1995). For strategic trade considerations to be relevant, a necessary condition is the existence of market power and thus ‘rents on the table’ that can be shifted from foreign companies to domestic firms. However, when we allow firms to exercise market power and compete in quantities, the resulting cost function estimates are very similar to our baseline estimates that assume price-taking behavior (see Appendix Table C.8). Given the fragmented nature of the industry where the largest firm accounts for less than 5% of the market, our results suggest that the shipbuilding industry is unlikely an effective target for strategic trade policies.

Given China’s prominent role in global trade, another reason to subsidize shipbuilding may be to boost its import and export sectors. Indeed, a larger worldwide fleet reduces transportation costs and thus increases trade; if Chinese exporters and importers face entry barriers or other frictions, these subsidies may be justifiable. To evaluate this argument, we would have to augment our model with a general equilibrium model of trade in order to predict changes in China’s exports and imports brought by the enlarged ship fleet due to the industrial policy; unfortunately this creates complications, such as computing the corresponding welfare gains, that fall beyond the scope of this paper.⁴⁷

Finally, it is worth noting that other considerations, including national security and military implications, as well as the desire to be the world leader in heavy industries (as stated in various government documents), might be equally relevant in motivating

⁴⁷In particular, to perform this calculation we need to compute the change in freight rates after the subsidies (Kalouptsidi 2018 calculates this at 3-5% in the case of Handysize vessels), the trade elasticity with respect to freight rates (Brancaccio *et al.* 2018 calculate this to be about 1), China’s total *maritime* trade (the exact portion of China’s trade that is carried by ships, which is unknown) and finally, the resulting gains from trade. It is the latter element that is complicated and requires overlaying our framework within a GE trade model.

these policies. Regardless of the motivation, our analysis provides cost estimates (and welfare losses) and the relative efficacy of implementing these policies that can be used as a guidance for future policies.

3.7 Conclusion

We empirically evaluate China's industrial policy, using the shipbuilding industry as a case study. While subsidies led to a significant increase in China's world market share and buttressed China's ascent into global influence, their long-run rate of returns are low and they entail substantial distortions. Counterfactual simulations indicate that the effectiveness of subsidies can be improved substantially when implemented counter-cyclically or targeted towards more productive firms. Our results shed light on the impact of industrial policies implemented in developing countries and reveal the importance of a proper policy design.

APPENDIX A
APPENDICES FOR CHAPTER 1

A.1 Monte Carlo Simulations

I describe Monte Carlo simulations of a simplified version of the model developed in Section 2.4. In the simplified version of the model, marginal costs are linear until production hits a binding capacity constraint. Trade costs are deterministic and only a function of the distance between countries. The purpose of the simulations is to illustrate the kinds of trade flows and prices generated by Cournot competition and contrast these with the trade flows and prices generated by competitive behavior.

A.1.1 Assumptions

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

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

I consider the following scenarios:

Table A.1: Monte Carlo Assumptions

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

1. S1, where firms have lots of capacity relative to demand. The numbers are chosen so that for every possible pair of demand shocks, firms will produce below capacity.
2. S2, where firms have very limited capacity. The numbers are chosen so that for every possible pair of demand shocks, firms will produce at full capacity even when they have no contracts.

I present some comparisons of the Cournot model and the perfectly competitive model. I look at both scenario S1 (where firms have lots of capacity) and scenario S2 (where firms have very limited capacity). For each scenario, I draw 1000 random demand shocks ε_1 and ε_2 from the uniform distribution and then solve for the perfectly competitive outcome and the Cournot outcome. Table A.2 shows average prices, trade flows, production and welfare across all 1000 demand shock realizations.

When firms have lots of capacity (scenario S1), firms withhold production under Cournot competitive, relative to the competitive allocation. This leads to higher prices and lower welfare under Cournot. Another feature of Cournot competition is that sellers are more likely to sell to faraway buyers. Under perfect competition, seller 1 only sells

to buyer 1 and seller 2 only to buyer 2. Under Cournot, seller 1 often sells to buyer 2 and seller 2 to buyer 1. This happens because the nearby seller withholds quantity from the nearest buyer, which sometimes makes it profitable for the faraway seller to enter the market.

When firms have very limited capacity (scenario S2), total production is pinned down by capacity and does not depend on the type of competition. However welfare is still lower under Cournot competition. There are two sources of this inefficiency. One is that transportation costs are not minimized under Cournot since buyers are more likely to buy from far away. For example seller 1 sells 83% of its output to the nearest buyer under perfect competition but only 62% under Cournot. This happens even in scenarios with similar demand in the two markets. An example is given by Table A.3 which shows that even when demand is very similar in the two markets, the Cournot outcome is worse from a welfare perspective because there is too much trade happening between faraway seller-buyer pairs.

A second source of inefficiency from Cournot is that when demand is higher in one market than other, oligopolists withhold quantity from the high-demand market and over-supply the low-demand market. This is shown in Table A.4, which looks at what happens when demand is very high in market 1 but very low in market 2. The competitive allocation involves selling all the gas to market 1; under Cournot, however, seller 2 withholds some of its output from market 1 and sells to market 2 instead, raising prices in market 1 and reducing them in market 2. As such, price differentials across regions are larger under Cournot than in perfect competition.

The Cournot model generates the following stylized facts. First, there is imperfect sorting of trade flows by distance: sellers often sell LNG to buyers that are very far away, despite the presence of nearby buyers. Second, sellers engage in price discrimination:

Table A.2: Comparison of competitive and Cournot outcome,

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

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

Table A.3: Scenario with moderate demand in both markets

ε_1	ε_2		Welfare	p_1	p_2	q_{11}	q_{12}	q_{21}	q_{22}
55.9	56.7	Competitive	1314	40.9	41.7	15	0	0	15
		Cournot	1270	41.2	41.5	9.4	5.6	5.4	9.6

Table A.4: Scenario with high demand in market 1, low demand in market 2

ε_1	ε_2		Welfare	p_1	p_2	q_{11}	q_{12}	q_{21}	q_{22}
78.3	41.0	Competitive	1689	48.3	41.0	15	0	15	0
		Cournot	1671	51.2	38.1	15	0	12.1	2.9

they earn higher markups from some buyers than other buyers. In the competitive model there is no price discrimination, and less likelihood that sellers will sell to faraway buyers. These are all consistent with LNG trade flows and prices that are observed in practice. It is also instructive to note the predictions which are similar for the Cournot and competitive models. Both models produce price dispersion and convergence: prices are different across regions during periods when demand is different, but prices are similar when demand shocks are similar (though the Cournot model predicts higher price differentials during periods with asymmetric demand). The Cournot model with capacity constraints also does not predict *average* prices that are higher than perfect

competition; thus the spatial nature of trade flows is essential in order to tell apart Cournot competition from perfect competition.

A.2 Derivation of first-order condition in the case of linear demand and quadratic costs

In the case of linear demand and quadratic costs, the first-order condition (1.10) becomes linear in the match cost v_{ij} , as we show below.

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

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

Because of the assumption of quadratic costs of production (see equation (1.13)), we can write expected costs and the derivative of expected costs with respect to S_{ij} as:

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

Plugging these expressions into equation (2.7), we get the following equation:

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

Collecting terms:

$$\begin{aligned} \beta(\varepsilon_j - Q_j^c - \sum_{i' \neq i} E(S_{i'j}(\mathbf{v}_{i'j}))) - (c_i + \delta(k_i)(q_i^c + \sum_{k \neq j} E(S_{ik}(\mathbf{v}_{ik}))) + d_{ij} + \phi(x_{ij})) - \\ (\beta(1 + \theta_i^c) + \delta(k_i))S_{ij}(\mathbf{v}_{ij}) - \mathbf{v}_{ij} \leq 0 \end{aligned} \quad (\text{A.1})$$

APPENDIX B
APPENDICES FOR CHAPTER 2

B.1 Long-term contracts and short-run allocations: Monte Carlo simulations

I describe Monte Carlo simulations of the model developed in Section 2.4. 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 welfare.

Assumptions

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 firms cannot produce beyond their capacity K_i and have constant marginal costs for any output below K_i . This is a special case of the hockey-stick cost function in equation (2.30) with $\delta_2 = \infty$ and $v = 1$. I set the marginal cost of production $\delta_1 = 4$. I assume that demand shocks in the two buying countries are uniformly and independently distributed: $\varepsilon_j \sim U[D_l, D_h]$, $j = 1, 2$.

I consider two different simulations (see Table B.2):

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

Table B.1: Assumptions maintained in all scenarios

Demand Slope	$\beta = 1$	
Demand Shock	$\varepsilon_j \sim U[d_l, d_h]$	
Demand Parameters	$D_l = 40$	$D_h = 80$
Marginal cost of production	$\delta_1 = 4$	
Shipping costs from seller 1	$d_{11} = 1$	$d_{12} = 5$
Shipping costs from seller 2	$d_{21} = 5$	$d_{22} = 1$

Table B.2: Simulations

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

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

For each of the two simulations, I consider a variety of possible contract configurations (Table B.3). I first look at the case where firms sign no contracts at all (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% of their contract quantity 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 spot output to the nearest buyer.

Table B.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
Sellers sign contracts	C2	1.25	0.75	0.75	1.25	2	2
	C3	4.37	2.63	2.63	4.37	7	7
	C4	8.75	5.25	5.25	8.75	14	14
	C5	9.375	5.625	5.625	9.375	15	15

Scenario 1 (S1): firms always produce below capacity

Table B.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 reducing the payoff they receive on their contracted output.

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

		Prices			Trade flows					
		Welfare			Seller 1		Seller 2		Q	K
	Contracts		p_1	p_2	q_{11}	q_{12}	q_{21}	q_{22}		
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.

Scenario 2 (S2): firms always produce at capacity

Table B.5 looks at the effect of contracts into the Cournot model when firms have limited capacity. Now contracts no longer affect total production, so the welfare gains identified by Allaz and Vila (1993) are no longer present. Moreover when the contracted quantities are sufficiently large, welfare decreases as we increase contracts. This shows that contracts are welfare-reducing when firms are fully capacity-constrained.

Table B.5: Effect of contracts when firms produce at capacity

		Prices			Trade flows					
		Welfare			Seller 1		Seller 2		Q	K
	Contracts		p_1	p_2	q_{11}	q_{12}	q_{21}	q_{22}		
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 demand shocks and asymmetric demand shocks in Table B.6 and Table B.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 are unable to do this. Table B.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. This is because sellers sell contracted LNG to buyer 2 even though buyer 1 is willing to pay a higher spot price than buyer 2. This also means that price differentials get larger because of contracts: comparing the second row with the last row, we can see that reducing contracted quantities would lead to smaller price differences across regions. (The results are similar when demand is low in market 1 but high in market 2). By contrast, as Table B.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 B.6: Effect of contracts when demand is high in market 1, low in market 2

		Prices			Trade flows					
		Welfare			Seller 1		Seller 2		Q	K
	Contracts		p_1	p_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 $(\varepsilon_1, \varepsilon_2)$. I solve for spot prices and allocation, taking any contracted flows as given. The table presents averages across realizations of ε where $\varepsilon_1 > 70$ and $\varepsilon_2 < 50$.

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

		Prices			Trade flows					
		Welfare			Seller 1		Seller 2		Q	K
	Contracts		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 $(\varepsilon_1, \varepsilon_2)$. I solve for spot prices and allocation, taking any contracted flows as given. The table presents averages across realizations of ε where $\varepsilon_1 > 70$ and $\varepsilon_2 > 70$.

APPENDIX C

APPENDICES FOR CHAPTER 3

This appendix discusses the calibration of the fixed production cost, details for the dynamic estimation (e.g., estimates of the first-stage policy functions and state-variable transitions), and implementation of the counterfactual analyses.

C.1 Estimation Details

C.1.1 Additional Charts and Tables

Table C.1: Shipbuilding National Industrial Policies

Year	Shipbuilding National Industrial Policies	Plan Period
2003	National Marine Economic Development Plan	2001-2010
2006	The 11th Five-Year Plan for National Economic and Social Development	2006-2010
2006	The Medium and Long Term Development Plan of Shipbuilding Industry	2006-2015
2007	The 11th Five-Year Plan for the Development of Shipbuilding Industry	2006-2010
2007	The 11th Five-Year Plan for the Development of Shipbuilding Technology	2006-2010
2007	11th Five-Year Plan for the Development of Ship Equipment Industry	2006-2010
2007	Guideline for Comprehensive Establishment of Modern Shipbuilding (2006-2010)	2006-2010
2007	Shipbuilding Operation Standards	2007-
2009	Plan on the Adjusting and Revitalizing the Shipbuilding Industry	2009-2011
2010	The 12th Five-Year Plan for National Economic and Social Development	2011-2015
2012	The 12th Five-Year Plan for the Development of the Shipbuilding Industry	2011-2015
2013	Plan on Accelerating Structural Adjustment and Promoting Transformation and Upgrading of the Shipbuilding Industry	2013-2015
2013	Shipbuilding Industry Standard and Conditions	2013-

Table C.2: Demand estimates

Dependent variable:	(1) Orders	(2) Orders	(3) Orders	(4) Orders
Price (bulk)	-2.34*** (0.77)	-1.67*** (0.64)	-2.07*** (0.69)	-2.12*** (0.75)
Price (tanker)	-2.66*** (0.60)	-1.46* (0.88)	-1.80** (0.78)	-1.76** (0.89)
Price (container)	-4.85*** (0.91)	-2.44*** (0.85)	-3.39*** (1.01)	-3.39*** (0.99)
Price*Post2006	1.34*** (0.18)	1.00*** (0.14)	1.15*** (0.15)	1.34** (0.55)
Backlog (log)	0.34 (0.25)	-1.00*** (0.33)	-0.78** (0.38)	-0.81** (0.37)
Freight rate (bulk)	2.84*** (0.45)	3.27*** (0.56)	3.35*** (0.57)	3.33*** (0.56)
Freight rate (tanker)	4.04*** (0.70)	3.24*** (0.68)	2.94*** (0.65)	2.91*** (0.65)
Freight rate (container)	6.45*** (0.87)	4.47*** (0.73)	4.69*** (0.77)	4.60*** (0.75)
US Wheat price		-0.12 (0.48)	-0.10 (0.48)	-0.12 (0.49)
Iron ore imports, China		2.62*** (0.90)	2.93*** (0.89)	3.01*** (0.92)
Middle East refinery production		1.37 (1.05)	1.84* (0.97)	1.66* (0.99)
World Car Trade		1.32*** (0.44)	2.08*** (0.49)	2.05*** (0.49)
Trend			-0.026** (0.011)	-0.020 (0.019)
Trend*Post2006				-0.0026 (0.0076)
R^2 , bulk	0.68	0.71	0.71	0.71
R^2 , tanker	0.26	0.33	0.35	0.36
R^2 , container	0.44	0.52	0.51	0.51

Note: N equals 64 for bulk and container and 61 for tankers. The freight rate is the Baltic Exchange Freight Index for bulk ships, Baltic Exchange Clean Tanker Index for tankers, and the Containership Timecharter Rate Index for containerships. The demand shifters include the US wheat price and total Chinese iron ore imports for bulk, Middle East refinery production for tanker, and world car trade for containership. We instrument ship prices using steel production and the steel ship plate price. Parameters are estimated using GMM.

Table C.3: Actual vs. Simulated Exit

	1999-2005	2006-2013	Total
Actual exits	5	43	48
Simulated exits	9	30	39

Note: We simulate the model 50 times from 1999 to 2013 under the baseline assumptions, and report the average number of exits across these simulations in the table above.

Table C.4: Actual vs. Simulated Entrants

	Pre	Post, Until 2008	Post, 2009+	Total
Actual entries	83	122	39	244
Simulated entries	69	125	38	232

Note: "Pre" refers to the period prior to 2004 for Zhejiang, and prior to 2006 for Jiangsu, Liaoning and Other regions. "Post, Until 2008" refers to the period between 2004 and 2008 for Zhejiang and between 2006 and 2008 for Jiangsu, Liaoning and Other regions. "Post, 2009+" refers to the period from 2009 onwards. We simulate the model 50 times from 1999 to 2013 under the baseline assumptions, and report the average number of entries across these simulations in the table above.

Table C.5: Impact of Subsidies on Ship Prices

	Bulk	Tanker	Container
Subsidies, 2006-08	16.3	20.0	17.2
No subsidies, 2006-08	17.6	21.2	17.7
% difference	8.2%	6.2%	3.1%
Subsidies, 2009-13	8.8	8.1	9.2
No Subsidies, 2009-13	10.2	9.0	9.5
% difference	16.5%	10.6%	3.7%

Note: Prices in 1000 RMB/CGT

Table C.6: Comparing Production and Investment Subsidies

	100% Production Subsidy	25% Production Subsidy	Investment Subsidy
Lifetime Revenue 2006-	2048	1765	1760
Lifetime Profit 2006-	768	610	602
Production Subsidy	219	42	0
Investment Subsidy	0	0	43
Δ Revenue /Subsidy	180%	229%	214%
Δ Net Profit/Subsidy	53%	78%	78%


Note: Each element in the table refers to the discounted sum from 2006 to 2050, averaged across all of the simulations. Δ Revenue/Subsidy and Δ Net Profit/Subsidy are defined as in Table 3.5. In the scenario “100% Production Subsidy”, we maintain the same production subsidy as in the baseline estimation, but shut down entry and investment subsidies. In the scenario “25% Production Subsidy”, we assume that the production subsidy per unit is 25% of the subsidy in the baseline estimates. In the scenario “Investment Subsidy”, we maintain the same investment subsidy as in the baseline estimation, but shut down entry and production subsidies. For each scenario, the government policy from 2014 onwards remains frozen at the 2013 policy.

Table C.7: Targeting subsidies to White List firms

	Subsidize All Firms After 2013	Subsidize White List firms After 2013	No Subsidies After 2013
Lifetime Revenue 2006-	2273	2230	2079
Lifetime Profit 2006-	831	820	702
Production Subsidy	249	213	126
Investment Subsidy	83	53	37
Entry Subsidy	302	301	303
Δ Revenue/Subsidy	115%	153%	
Δ Net Profit/Subsidy	34%	85%	


Note: Each element in the table refers to the discounted sum from 2006 to 2050, averaged across all simulations. For example, “Lifetime Profit 2006” refers to the discounted sum (in 2006) of profits earned by firms from 2006 to 2050. Δ Revenue/Subsidy and Δ Net Profit/Subsidy are measures of the rate of return on subsidies, relative to column (3), where the industry is subsidized until 2013 and then all subsidies are discontinued.

Figure C.1: China's Expansion in Major Industries




Source: NBS. The output of the auto, auto parts and steel industries are plotted on the left vertical axis, while the output of the solar industry is plotted on the right vertical axis.

Figure C.2: Simulated vs. actual investment




Note: For the simulated investment employed in the model, we use the estimated parameters and value function, randomly draw v for every observation, and calculate optimal investment.

Figure C.3: HHI for Chinese Shipbuilding, with and without subsidies



Notes: The HHI reported in the above figure is calculated using all Chinese yards in any given year. It is thus a measure of concentration within the Chinese shipbuilding industry.

Figure C.4: Industry Profits Under Different White Lists



Note: In all scenarios, we pick a set of firms in the White List in 2013 and force all other firms to exit. In the scenario “Subsidies and Actual White List”, the government subsidizes the industry until 2013 and then keeps the observed White List. In the scenario “Subsidies and Optimal White List”, the government subsidizes the industry until 2013 and then chooses a White List based on observed profitability. In the scenario “No Subsidies and Optimal White List”, the industry does not receive any subsidies, and in 2013 a White List of firms is selected based on observed profitability.

C.1.2 Calibrating the Fixed Cost

The NBS data include information on operating costs, which allows us to calibrate the fixed cost of production. A firm's total production cost is equal to:

$$C_{jt} = c_0 + C(q_{jt})$$

where $C(q_{jt})$ is the variable cost of taking q_{jt} orders that is estimated from the Clarkson data, as discussed in Section 3.4.1.

Let C_{jt}^{NBS} denote the accounting operating costs, which include the costs of both ship production and ship repairs, a common practice in this industry. We follow the standard assumption in the production literature that the cost share of ship production is the same as its revenue share and obtain the accounting operating cost of ship production as:

$$\hat{C}_{jt} = C_{jt}^{NBS} * (R_j^{Clarkson} / R_j^{NBS})$$

where $R_j^{Clarkson} = \sum_t R_j^{Clarkson}$ denotes j 's lifetime revenue from building new ships that is reported in Clarkson and $R_j^{NBS} = \sum_t R_{jt}^{NBS}$ denotes its lifetime revenue in NBS.

We use two approaches to estimate the fixed cost c_0 ; both deliver similar results. The first approach uses the quarters with zero production (so that the variable production cost is zero) and the accounting costs $\hat{C}_{jt}$ (after adjusting for repairs) in the same periods to infer the fixed cost. The second approach uses the difference between a shipyard's average operating costs and the average estimated variable cost of production:

$$c_0 = \frac{1}{T} \sum_t [\hat{C}_{jt} - C(q_{jt})]$$

Note that the calibrated fixed cost of operation is the same for all firms. While in principle we could allow the fixed cost to vary by firm characteristics, our data is not rich enough to deliver a precise estimate.

C.1.3 Production cost estimates: robustness exercises

In this section, we carry out various exercises to explore how robust our estimates of production cost are to various modelling assumptions.

First, we relax the assumption of price-taking behavior by firms, and estimate production costs assuming the firms play a Cournot game among each other. Table C.8 shows that the estimated coefficients remain quantitatively similar. This is not unexpected since the shipbuilding industry is unconcentrated, with the HHI ranging from 230 to 720 during the sample period.

Next, we look at how our estimates of ship production cost change when we pool together data from China, Japan and South Korea. The challenge with doing so is that the capital stock is unobserved for firms in Japan and South Korea. We set the capital variable to zero for Japanese and South Korean yards and use country dummies to control for these capital costs. The results are illustrated in Table C.9. The key coefficients are generally very similar. The subsidy is estimated to be significantly higher in the 2006-08 period but somewhat lower in the 2009+ period.

In Table C.11, we report results from other robustness exercises. First, we allow for a time trend in the cost function. This captures, for example, changes in the technology of shipbuilding over time. The time trend is estimated to be very small in magnitude and including the trend has little effect on our estimated cost parameters. Secondly, we repeat the analysis on a sub-sample where we drop any Chinese yards that entered after the policies were announced. The results are broadly similar with the estimated production subsidy lower in the 2006-08 period relative to the baseline specification.

Finally, we examine evidence of learning-by-doing by shipyards. First, we evaluate within-firm learning-by-doing by allowing a firm's marginal cost to depend on its

cumulative past production. As shown in Table C.10, a larger past production leads to higher marginal costs, which is inconsistent with there being any within-firm learning-by-doing. Second, we allow a firm's marginal cost to depend on the industry cumulative output, as a crude test of industry-wide learning-by-doing (where firms learn from each other). Without instrumenting for the industry cumulative output, this exercise is likely to over-estimate spillover effects: if there are common unobserved shocks that raise the output of all firms, it will look as though there are positive spillover effects. Despite this, we find limited evidence for spillover effects, as we can see in the third panel of Table C.10. Marginal costs increase with the cumulative industry production for tanker and container and only modestly decrease with the cumulative industry production for bulk, though none of these coefficients is statistically significant.

Table C.8: Cost function estimates: perfect competition vs. Cournot competition

	Bulk		Tanker		Container	
	Coefficient	T-stat	Coefficient	T-stat	Coefficient	T-stat
Perfect competition						
Capital (bill RMB)	-2.43	-2.96	-2.61	-1.80	-2.19	-2.01
Backlog	-1.56	-5.29	-4.44	-5.04	-2.88	-3.34
2006-2008	-1.51	-2.62				
2009+	-1.38	-2.37				
Cournot						
Capital (bill RMB)	-2.67	-2.85	-2.89	-1.74	-2.44	-1.93
Backlog	-1.80	-5.03	-5.02	-4.97	-3.30	-3.19
2006-2008	-2.10	-3.01	-2.10	-3.01	-2.10	-3.01
2009+	-1.22	-1.78	-1.22	-1.78	-1.22	-1.78
N	4886		4977		2504	

Note: The first panel reports the baseline cost function estimates from Table 3.2, which assume perfect competition. The second panel reports the cost estimates assuming that firms compete in quantities. The table reports estimates for key coefficients. Full tables are available upon request. Standard errors bootstrapped using 500 bootstrap samples.

Table C.9: Cost function estimates, pooling data across China/Japan/South Korea

	Bulk		Tanker		Container	
	Coefficient	T-stat	Coefficient	T-stat	Coefficient	T-stat
Chinese yards						
Capital (bill RMB)	-2.43	-2.96	-2.61	-1.80	-2.19	-2.01
Backlog	-1.56	-5.29	-4.44	-5.04	-2.88	-3.34
China 2006-2008	-1.51	-2.62				
China 2009+	-1.38	-2.37				
N	4886		4977		2504	
Chinese/Japanese/Korean yards						
Capital (bill RMB)	-2.92	-3.06	-2.06	-1.55	-1.41	-1.33
Backlog	-2.09	-6.63	-4.50	-6.38	-3.06	-4.40
China 2006-2008	-2.79	-4.57				
China 2009+	-0.90	-1.56				
N	10013		10429		4661	

Note: The first panel reports the baseline cost function estimates from Table 3.2, which only uses data on Chinese yards. In the second panel, we pool together Chinese/Japanese/Korean yards. To account for missing capital stock data for non-Chinese yards, we set the capital variable to zero for Japanese and Korean yards, and allow the constant term to differ by country. We also allow the backlog coefficients to differ by country. Backlog coefficients for Japan and Korea are not reported above. Standard errors bootstrapped using 500 bootstrap samples.

Table C.10: Cost function estimates: learning

	Bulk		Tanker		Container	
Type-specific	Coefficient	T-stat	Coefficient	T-stat	Coefficient	T-stat
Baseline specification						
Capital (bill RMB)	-2.92	-3.06	-2.06	-1.55	-1.41	-1.33
Backlog	-2.09	-6.63	-4.50	-6.38	-3.06	-4.40
Allow for within-firm learning						
Capital (bill RMB)	-1.94	-2.12	-1.95	-1.52	-1.11	-1.22
Backlog	-1.45	-5.05	-4.41	-5.06	-0.90	-1.30
Cumulative Q	0.07	4.72	0.09	5.59	0.01	4.00
Allow for within-firm and industry-wide learning						
Capital (bill RMB)	-2.04	-2.16	-3.41	-1.69	-2.04	-1.32
Backlog	-1.35	-4.58	-6.60	-4.02	-1.58	-1.16
Cumulative Q	0.08	4.76	0.13	4.39	0.02	2.94
Cumulative Q, China	-0.03	-1.60	0.18	1.47	0.33	1.30

Note: The first panel repeats key coefficients from the second specification reported in Table C.9. The second panel includes all regressors from the first panel, as well as each firm's cumulative past production. The third panel includes all regressors from the the first panel, each firm's cumulative past production, and the country's cumulative past production. Standard errors bootstrapped using 500 bootstrap samples. We pool together Chinese/Japanese/Korean yards.

Table C.11: Cost function estimates: additional robustness checks

	Bulk		Tanker		Container	
	Coefficient	T-stat	Coefficient	T-stat	Coefficient	T-stat
Baseline specification						
Capital (bill RMB)	-2.92	-3.06	-2.06	-1.55	-1.41	-1.33
Backlog	-2.09	-6.63	-4.50	-6.38	-3.06	-4.40
China 2006-2008	-2.79	-4.57				
China 2009+	-0.90	-1.56				
Add time trend						
Capital (bill RMB)	-2.91	-3.11	-2.06	-1.56	-1.41	-1.33
Backlog	-2.09	-6.41	-4.50	-6.42	-3.06	-4.40
China 2006-2008	-2.78	-4.19				
China 2009+	-0.89	-1.35				
Trend	-0.02	-0.03				
Existing yards						
Capital (bill RMB)	-3.33	-3.18	-2.49	-1.41	-0.44	-0.40
Backlog	-3.16	-6.38	-5.15	-6.13	-3.66	-4.25
China 2006-2008	-2.08	-2.72				
China 2009+	-1.01	-1.25				

Note: The first panel repeats key coefficients from the second specification reported in Table C.9. The second panel includes all regressors from the first panel, as well as a quarterly time trend. The third panel repeats the regression from the first panel on a sub-sample of shipyards where we exclude any Chinese yards that entered after the policies were announced. Standard errors bootstrapped using 500 bootstrap samples. We pool together Chinese/Japanese/Korean yards.

C.1.4 Estimating the Investment Policy Function

Our baseline estimator of the investment policy function does not account for the fact that investment is non-negative. To address this issue, we perform two robustness checks. The first is a Tobit model that assumes $h_2(v_{jt})$ is normally distributed. The second approach assumes that the median of $h_2(v_{jt})$ is zero and estimates $h_1(s_{jt})$ using the Censored Least Absolute Deviation estimator (CLAD) that was first proposed by Powell (1984) and later extended by Chernozhukov and Hong (2002). In this note, we describe how we implement the second approach based on the CLAD estimator.

The investment policy function is assumed to be additive in the observed state variables and the unobserved investment cost shock:

$$I_{jt}^* = h_1(s_{jt}) + h_2(v_{jt})$$

$$I_{jt} = \max(I_{jt}^*, 0)$$

where the second equation makes it explicit that investment is non-negative. Powell (1984) showed that we can recover $h_1(s)$ through the Censored Least Absolute Deviations estimator (CLAD) while normalizing the median of $h_2(v_{jt})$ to 0. Once we obtain the CLAD estimate $\hat{h}_1(s)$, we treat $I_{jt} - \hat{h}_1(s_{jt})$ as data with the goal of estimating $h_2(v_{jt})$ with the truncated data:

$$\tilde{i}_{jt} \equiv I_{jt} - \hat{h}_1(s_{jt}) = \max(h_2(v_{jt}), -\hat{h}_1(s_{jt})), \text{ or}$$

$$\tilde{i}_{jt} = \max(h_2(v_{jt}), \bar{h}_{jt})$$

where in the second equation we use $\bar{h}_{jt}$ to denote $-\hat{h}_1(s_{jt})$.

Note that the level of truncation $\bar{h}_{jt}$ varies across observations. We use the observed probability of truncation (zero or negative investment) to back out the level of the investment shock that induces truncation, conditioning on the observed state variables

(let Φ denote the CDF of a standard normal):

$$\begin{aligned} Pr(\tilde{i}_{jt} > \bar{h}_{jt} | \bar{h}_{jt}) &= Pr(h_2(v_{jt}) > \bar{h}_{jt}) = Pr(v_{jt} < h_2^{-1}(\bar{h}_{jt})) = Pr(v_{jt} < \bar{v}_{jt}) \\ &= \Phi(\bar{v}_{jt}), \text{ or} \\ \bar{v}_{jt} &= \Phi^{-1}(Pr(\tilde{i}_{jt} > \bar{h}_{jt} | \bar{h}_{jt})) \end{aligned}$$

where $Pr(\tilde{i}_{jt} > \bar{h}_{jt} | \bar{h}_{jt})$ can be estimated either via kernel methods, or by approximating the cutoff value $\bar{v}(\bar{h}_{jt})$ by a flexible function of $\bar{h}_{jt}$ and carrying out a probit regression.

To estimate $h_2(v_{jt})$, we categorize all the uncensored observations (where $\tilde{i}_{jt} > \bar{h}_{jt}$) into distinct bins. Specifically, suppose the thresholds are $\{\bar{h}_1, \bar{h}_2, \dots, \bar{h}_{R+1}\}$. Then any uncensored observation $\tilde{i} \in (\bar{h}_r, \bar{h}_{r+1}]$ is placed in bin r . We carry out the BBL inversion separately for each bin. In particular, if $i^* = \max(h_2(v^*), \bar{h}_{jt})$ for some arbitrary v^* , where i^* lies in bin r , then the following expression must hold:

$$\begin{aligned} F(i^* | i^* \in (\bar{h}_r, \bar{h}_{r+1}]) &= Pr(\tilde{i} \leq i^* | i^* \in (\bar{h}_r, \bar{h}_{r+1})) \\ &= Pr(v \geq v^* | \bar{v}_{r+1} < v < \bar{v}_r) \\ &= \frac{\Phi(\bar{v}_r) - \Phi(v^*)}{\Phi(\bar{v}_r) - \Phi(\bar{v}_{r+1})} \end{aligned}$$

In other words,

$$i^* = F^{-1}\left(\frac{\Phi(\bar{v}_r) - \Phi(v^*)}{\Phi(\bar{v}_r) - \Phi(\bar{v}_{r+1})}\right) \text{ for } \bar{v}_{r+1} < v^* < \bar{v}_r$$

It is easy to verify that this estimator nests the uncensored example as a special case and allows us to better address censoring by increasing the number of bins. Monte Carlo simulations suggest that a small number of bins (say five) can lead to surprisingly well-behaved estimates with minimal bias in the estimated function $h_2(v)$.

C.1.5 First-stage Policy Functions and State Transition Estimates

This section presents the first-stage estimates of the investment and exit policy functions, as well as the state transition process. Table C.12 presents estimates of the transition process for the prices of bulkers, tankers, containerships, and steel. Table C.13 reports the estimated exit policy function. Table C.14 reports the estimated investment policy function using OLS, Tobit, and CLAD.

Table C.12: AR(1) estimates for state transition processes

	Bulk	Tanker	Container	Steel
Constant	0.88 (0.87)	0.70 (0.94)	1.25 (1.11)	-0.023 (0.37)
Post	3.44 (2.28)	3.63 (3.43)	1.80 (3.33)	2.32 (1.01)
Price (t-1)*Pre	0.86 (0.12)	0.92 (0.086)	0.88 (0.090)	0.89 (0.19)
Price (t-1)*Post	0.86 (0.072)	0.86 (0.095)	0.88 (0.10)	0.69 (0.11)
Trend*Pre	0.042 (0.033)	0.038 (0.028)	0.029 (0.024)	0.024 (0.027)
Trend*Post	-0.058 (0.027)	-0.054 (0.041)	-0.040 (0.040)	-0.022 (0.013)
N	57	57	57	57
R ²	0.95	0.97	0.96	0.80

Note: The dependent variable is the price in quarter t , or $p(t)$. Standard errors in parenthesis. "Pre" refers to 2005Q4 or earlier. "Post" refers to 2006Q1 or later. The sample ranges from 1999 Q4 to 2013Q4.

Table C.13: Estimates of exit policy function

	(1)		(2)	
	Coefficient	SE	Coefficient	SE
Constant	-0.57	(0.97)	-0.56	(1.02)
K	0.05	(0.35)	0.54	(0.43)
K^2	-0.05	(0.12)	-0.16	(0.15)
2006-2008	-0.57	(0.41)	-0.64	(0.43)
2009+	-0.47	(0.41)	-0.72	(0.44)
$\bar{s}_{jt}$	-0.01	(0.01)	-0.04	(0.02)
Bulk price	0.36	(0.12)	0.36	(0.12)
Tanker price	-0.18	(0.11)	-0.16	(0.11)
Container price	-0.22	(0.10)	-0.25	(0.11)
Steel price	-0.06	(0.07)	-0.10	(0.08)
Jiangsu			0.77	(0.24)
Zhejiang			0.58	(0.19)
Liaoning			1.04	(0.28)
N	4605		4605	
Log-likelihood	-239.30		-229.74	
Pseudo-R2	0.09		0.12	

Note: We carry out a probit regression of a binary indicator of exit on basis functions of the states. We restrict the estimation to 1999-2011, because firm exits in 2012 and 2013 are not reliably measured.

Table C.14: Estimates of the investment policy function

	(1) OLS	(2) Tobit	(3) CLAD
Constant	-0.066 (7.54)	-12.2 (8.17)	-31.9*** (4.09)
B-spline 1 Capital	-69.7*** (22.0)	-63.8*** (17.2)	-69.6*** (1.67)
B-spline 2 Capital	-74.7*** (17.7)	-71.7*** (13.5)	-68.2*** (1.41)
2006-08	6.42*** (1.60)	4.59** (2.32)	17.9*** (0.74)
2009+	2.70 (2.20)	3.79 (3.03)	3.55** (1.80)
$\bar{s}_{jt}$	0.74*** (0.11)	0.87*** (0.087)	1.44*** (0.040)
Bulk price	2.05*** (0.46)	1.97*** (0.57)	1.34*** (0.30)
Tanker price	0.48 (0.93)	1.89* (1.14)	0.81*** (0.13)
Container price	-1.25 (0.87)	-1.49 (1.06)	-0.55 (0.34)
Steel price	-2.49*** (0.53)	-4.44*** (0.61)	-4.38*** (0.19)
N	4286		
$N(I > 0)$	3301		
$N(I = 0)$	985		

Note: In Column (1), we carry out an OLS regression of investment (I) on basis functions of the states, including both observations with $I > 0$ and $I = 0$. In (2), we estimate the policy function using a Tobit regression of I on the basis functions. In (3), we estimate the investment policy function using a censored least absolute deviations estimator. $\bar{s}_{jt}$ is an index capturing the effect of backlog, age, ownership, region, and size on a firm's per-period payoffs. Investment is measured in million RMBs in these regressions.

C.1.6 Value function approximation

As discussed in the main text, we approximate the value function $V(s_{jt})$ via B-spline basis functions $V(s_{jt}) = \sum_{l=1}^L \gamma_l^0 u_l(s_{jt})$ and impose the Bellman equation as a constraint when estimating the parameter governing investment costs and scrap value. We now discuss some additional issues related to how we approximate the value functions.

Constructing Basis Functions In our model of the Chinese shipbuilding industry, firm value functions are in principle a function of a large number of state variables. However, several state variables enter the shipyard's payoff as a single index $s_{jmt} \beta_{sm}$ in the marginal cost of production (3.10), including the shipyard's region, ownership, size, age, and backlog. As such, instead of keeping track of each state individually, we collapse them into a single-dimensional state using the estimated coefficients:

$$\bar{s}_{jt} = - \sum_m s_{jmt} \beta_{sm} \quad (\text{C.1})$$

We use $\bar{s}_{jt}$ as a measure of a firm's observed cost efficiency: a higher $\bar{s}_{jt}$ is associated with a lower marginal cost and a higher variable profit. Our approach of collapsing firm-level state variables into a single index is similar in spirit to Hendel and Nevo (2006) and Nevo and Rossi (2008) that use the “inclusive value” to capture the impact of changing product attributes on future profits. We further assume that $\bar{s}_{jt}$ evolves via a simple rule $\bar{s}_{jt+1} = \alpha_0 + \alpha_1 \bar{s}_{jt}$, which almost perfectly forecasts $\bar{s}_{jt+1}$ in period t since all but one of the variables in $\bar{s}_{jt}$ are deterministic.

Therefore, the state variables in the dynamic estimation are the capital stock, the price for each ship type, the steel price, and $\bar{s}_{jt}$ (which subsumes the remaining firm characteristics); as well as two policy dummies for the periods 2006-08 and post 2009, respectively. The basis functions are flexible third order B-splines (i.e. quadratic

piecewise polynomials). Given our focus on investment, we use two knots (and have experimented with more knots) in forming the B-splines for capital. The total number of basis functions is 44.

Estimating approximating coefficients We search for $\{\gamma_l\}_{l=1}^L$ that minimize the violation of the Bellman equation (3.12) given the dynamic parameters:

$$\{\gamma_l\}_{l=1}^L = \arg \min_{\gamma} \|V(s_{jt}; \gamma) - \pi(s_{jt}) - \hat{p}^x(s_{jt})\sigma_\phi - CV(s_{jt}; \gamma)\|_2 \quad (\text{C.2})$$

where $\hat{p}^x(s_{jt})$ and $\hat{i}^*(s_{jt}, \mathbf{v}_{jt})$ are the estimated first-stage exit and investment policy functions, respectively, $CV(s_{jt}; \gamma) = E_{\mathbf{v}_{jt}}\{-C^i(\hat{i}^*(s_{jt}, \mathbf{v}_{jt})) + \beta E[V(s_{jt+1}; \gamma)|s_{jt}, \hat{i}_{jt}^*]\}$ is the continuation value evaluated at these estimated policy functions, and $\|\cdot\|_2$ is the L^2 norm.

Equation (C.2) is imposed as a constraint in the estimation of the dynamic parameters. Specifically, for each guess of the dynamic parameters, we solve for the $\{\gamma_l\}_{l=1}^L$ satisfying equation (C.2), and use these estimated approximating coefficients to construct the sample log likelihood in equation (3.13)

Recovering the approximating coefficients γ requires specifying the set of state values on which to evaluate the Bellman constraint. We construct a sample that ensures sufficient variation in each of the state variables. First, we include all the N states observed in the sample. Second, we randomly draw N_{add} additional states to span the full range of the state variables. The coefficients γ are recovered using these $N + N_{add}$ states. This approach is similar to Sweeting (2013).

These additional states are instrumental in getting a good approximation of the value function, for two reasons. First, some states (for example, ship prices and the steel price) are highly correlated in the data, which makes it challenging to separately identify

the coefficients on basis functions formed from these state variables if we only use the observed states. Second, some regions of the state space have a limited number of observations. Both of these problems can be mitigated by adding randomly drawn states, which avoids multicollinearity between states and ensures sufficient data points across all regions of the state space.

C.2 Implementation of Counterfactual Analyses

Each of the counterfactual experiments involves two steps: first, solving for the new Bellman equation and policy functions, and second, simulating the industry forward until 2099. Here we briefly explain how to implement the first step through a fixed point algorithm:

1. Compute expected profits $\pi(s)$ at all states.
2. Start with an initial guess of the exit policy function $p^{0,x}(s)$ and investment policy function $i^0(s, v)$.
3. Update the policy functions. At each iteration j :
 - Solve for the value function coefficients γ^{j+1} using the equation $V^{j+1}(s) = \pi(s) + p^{j,x}\sigma + CV^{j+1}(s)$.
 - Update the investment policy function to $i^{j+1}(s, v)$ by solving the investment FOC, using V^{j+1} and CV^{j+1} . As the value function is approximated by cubic B-splines, the investment policy function has an analytic solution.
 - Update the exit policy function to $p^{j+1,x}$ using V^{j+1} and CV^{j+1} .
 - Check whether $\|p^{j+1,x}(s) - p^{j,x}(s)\| < tol$ and $\|i^{j+1}(s, v) - i^j(s, v)\| < tol$, where tol is a pre-assigned tolerance level.

C.3 A Simple Model of Subsidies

This is a static model with homogeneous firms. Each firm has a starting capital stock of K_0 . Price is equal to P . Marginal cost of production equals $MC(q_t) = \alpha - \beta K + \delta q_t$. Total cost of investment equals $C_I(I) = c_1 I + (c_2/2)I^2$. The firm chooses q and I simultaneously to maximize profits:

$$V(K_0) = \max_{q,I} Pq - \left((\alpha - \beta(K_0 + I))q - \frac{\delta}{2}q^2 \right) - \left(c_1 I + \frac{c_2}{2}I^2 \right)$$

The optimal quantity and investment are denoted by q^* and I^* , respectively.

Now suppose that the government introduces production subsidies of τ_p per unit. For simplicity, we assume that the firm only adjusts its level of production and not investment; thus investment remains fixed at I^* . The new level of production, $\hat{q}$, is:

$$\hat{q} = q^* + \frac{\tau_p}{\delta}$$

Alternatively suppose the government introduces investment subsidies of τ_i per unit. The new level of investment, $\hat{I}$, is:

$$\hat{I} = I^* + \frac{\tau_i}{c_2}$$

Below we provide expressions for the return to subsidies, which is the change in industry profits from the subsidies divided by the cost to the government of providing the subsidies, as well as the deadweight loss from subsidies:

$$\text{DWL from Prod. Subsidies} = \tau_p^2 / 2\delta$$

$$\text{DWL from Invest Subsidies} = \tau_i^2 / 2c_2$$

$$\text{Return to Prod. Subsidies} = (q^* + \frac{\tau_p}{\delta}) / (q^* + \frac{\tau_p}{\delta})$$

$$\text{Return to Invest Subsidies} = (I^* + \frac{\tau_i}{c_2}) / (I^* + \frac{\tau_i}{c_2})$$

Holding the adjustment cost parameters c_2 and δ fixed, the return to subsidies is increasing in I^* (for investment) and q^* (for production). In other words, subsidizing “better” firms leads to higher returns.

Our derivation of the DWL shows that the magnitude of DWL is independent of whether a firm has low or high marginal costs: τ_p^2/δ . Essentially, all firms (large or small) increase their output by the same amount when they receive the same subsidy. However, the return to subsidies is higher for low-cost firms than high-cost ones. This is because low-cost firms receive a higher absolute amount of subsidies due to the fact that they produce a higher quantity. Thus the DWL loss is divided by a larger denominator which means the per-dollar return to subsidies is higher.

BIBLIOGRAPHY

- ABEL, A. B. and EBERLY, J. C. (1994). A unified model of investment under uncertainty. *American Economic Review*, **84** (5), 1369–1384.
- ACKERBERG, D., BENKARD, C. L., BERRY, S. and PAKES, A. (2007). Econometric tools for analyzing market outcomes. In J. Heckman and E. Leamer (eds.), *The Handbook of Econometrics. Vol. 6A.*, Amsterdam: North-Holland.
- AGERTON, M. (2017). Global lng pricing terms and revisions: An empirical analysis. *The Energy Journal*, **38** (1).
- AGHION, P., HEMOUS, D. and KHARROUBI, E. (2014). Cyclical fiscal policy, credit constraints, and industry growth. *Journal of Monetary Economics*, **62**, 41–58.
- , JING, C., DEWATRIPONT, M., LUOSHA, D., HARRISON, A. and LEGROS, P. (2015). Industrial policy and competition. *American Economic Journal: Macroeconomics*, **7** (4), 1–32.
- AGUIRREGABIRIA, V. and NEVO, A. (2013). Recent developments in empirical io: Dynamic demand and dynamic games. In D. Acemoglu, M. Arellano and E. Deckel (eds.), *Advances in Economics and Econometrics. Vol. 3, Econometrics*.
- and SUZUKI, J. (2014). Identification and counterfactuals in dynamic models of market entry and exit. *Quantitative Marketing and Economics*, **12** (3), 267–304.
- ALDY, J., GERARDEN, T. D. and SWEENEY, R. L. (2018). Investment versus output subsidies: Implications of alternative incentives for wind energy, nBER Working Paper No. 24378.
- ALLAZ, B. and VILA, J.-L. (1993). Cournot competition, forward markets and efficiency. *Journal of Economic Theory*, **59** (1), 1–16.

- AMSDEN, A. H. (1989). *Asia's Next Giant: South Korea and Late Industrialization*. Oxford University Press.
- ANDERSON, J. E. and VAN WINCOOP, E. (2004). Trade costs. *Journal of Economic literature*, **42** (3), 691–751.
- ASKER, J., COLLARD-WEXLER, A. and DE LOCKER, J. (2019). (mis)allocation, market power, and global oil extraction. *American Economic Review*, **109** (4), 1568–1615.
- AW, B. Y., ROBERTS, M. J. and XU, D. Y. (2011). R&d investment, exporting, and productivity dynamics. *American Economic Review*, **101** (4), 1312–1344.
- BAI, J. and LIU, J. (2019). The impact of intranational trade barriers on exports: Evidence from a nationwide vat rebate reform in china.
- BAJARI, P., BENKARD, L. and LEVIN, J. (2007). Estimating dynamic models of imperfect competition. *Econometrica*, **75** (5), 1331–1370.
- BALDWIN, R. E. (1969). The case against infant-industry tariff protection. *Journal of Political Economy*, **77** (3), 295–305.
- and KRUGMAN, P. R. (1988). Market access and international competition: A simulation study of 16k random access memories. In R. C. Feenstra (ed.), *Empirical Methods for International Trade*, Cambridge, Mass.: MIT Press.
- BALLARD, C. L., SHOVEN, J. B. and WHALLEY, J. (1985). General equilibrium computations of the marginal welfare costs of taxes in the united states. *American Economic Review*, **75** (1), 128–138.
- BARRO, R. (1979). On the determination of public debt. *Journal of Political Economy*, **87**, 940–971.

- BARSEGHYAN, L., BATTAGLINI, M. and COATE, S. (2013). Fiscal policy over the real business cycle: A positive theory. *Journal of Economic Theory*, **148**, 2223–2265.
- BARWICK, P. J., CAO, S. and LI, S. (2020). Local protectionism, market structure, and social welfare: China’s automobile market.
- and PATHAK, P. (2015). The cost of free entry: An empirical analysis of real estate agents in greater boston. *RAND Journal of Economics*, **46** (1), 103–145.
- BENKARD, C. L. (2004). A dynamic analysis of the market for wide-bodied commercial aircraft. *The Review of Economic Studies*, **71** (3), 581–611.
- BLOOM, N., GRIFFITH, R. and VAN REENEN, J. (2002). Do r&d tax credits work? evidence from a panel of countries 1979-1997. *Journal of Public Economics*, **85** (1), 1–31.
- BOIK, A. (2017). The empirical effects of competition on third-degree price discrimination in the presence of arbitrage. *Canadian Journal of Economics*, **50** (4), 1023–1036.
- BORENSTEIN, S., BUSHNELL, J. B. and WOLAK, F. A. (2002). Measuring market inefficiencies in california’s restructured wholesale electricity market. *American Economic Review*, **92** (5), 1376–1405.
- BRANCACCIO, G., KALOUPSTIDI, M. and PAPAGEORGIOU, T. (2018). Geography, search frictions and endogenous trade costs, working Paper.
- BRANDER, J. E. (1995). Strategic trade policy. In G. Grossman and K. Rogoff (eds.), *Handbook of International Economics, Chapter 27, Vol. III*.
- BRANDT, L., BIESEBROECK, J. V. and ZHANG, Y. (2012). Creative accounting

- or creative destruction? firm-level productivity growth in chinese manufacturing. *Journal of Development Economics*, **97**, 339–351.
- BRESNAHAN, T. F. (1982). The oligopoly solution concept is identified. *Economics Letters*, **10** (1), 87–92.
- (1987). Competition and collusion in the american automobile oligopoly: The 1955 price war. *Journal of Industrial Economics*, **35** (4), 457–482.
- BRICKLEY, J. A., MISRA, S. and HORN, R. L. V. (2006). Contract duration: Evidence from franchising. *The Journal of Law and Economics*, **49** (1), 173–196.
- BRITO, D. and HARTLEY, P. (2007). Expectations and the evolving world gas market. *The Energy Journal*, **28** (1).
- BRUCE, N. (1990). Measuring industrial subsidies.
- BUSHNELL, J. B., MANSUR, E. T. and SARAVIA, C. (2008). Vertical arrangements, market structure, and competition: An analysis of restructured us electricity markets. *American Economic Review*, **98** (1), 237–266.
- CHEN, Z., JIANG, X., LIU, Z., SUÁREZ SERRATO, J. C. and XU, D. (2019a). Tax policy and lumpy investment behavior: Evidence from china’s vat reform.
- , LIU, Z., SERRATO, J. C. S. and XU, D. Y. (2019b). Notching r&d investment with corporate income tax cuts in china.
- CHERNOZHUKOV, V. and HONG, H. (2002). Three-step censored quantile regression and extramarital affairs. *Journal of the American Statistical Association*, **97** (459), 872–882.

- CHIPTY, T. and SNYDER, C. M. (1999). The role of firm size in bilateral bargaining: A study of the cable television industry. *Review of Economics and Statistics*, **81** (2), 326–340.
- COLLARD-WEXLER, A. (2013). Demand fluctuations in the ready-mix concrete industry. *Econometrica*, **81** (3), 1003–1037.
- COLLINS, G. and GRUBB, M. C. (2008). A comprehensive survey of china’s dynamic shipbuilding industry, china Maritime Studies Institute. U.S. Naval War College.
- COOPER, R. W. and HALTIWANGER, J. C. (2006). On the nature of capital adjustment costs. *The Review of Economic Studies*, **73** (3), 611–633.
- CRAWFORD, G. S. and YURUKOGLU, A. (2012). The welfare effects of bundling in multichannel television markets. *American Economic Review*, **102** (2), 643–85.
- CRISCUOLO, C., MARTIN, R., OVERMAN, H. G. and VAN REENEN, J. (2019). Some causal effects of an industrial policy. *American Economic Review*, **109** (1), 48–85.
- CROCKER, K. J. and MASTEN, S. E. (1988). Mitigating contractual hazards: Unilateral options and contract length. *RAND Journal of Economics*, pp. 327–343.
- DAFNY, L. S. (2010). Are health insurance markets competitive? *American Economic Review*, **100** (4), 1399–1431.
- DAS, S., ROBERTS, M. J. and TYBOUT, J. R. (2007). Market entry costs, producer heterogeneity, and export dynamics. *Econometrica*, **75** (3), 837–873.
- DIXIT, A. (1984). International trade policy for oligopolistic industries. *The Economic Journal*, **94**, 1–16.
- DORASZELSKI, U., LEWIS, G. and PAKES, A. (2018). Just starting out: Learning and equilibrium in a new market. *American Economic Review*, **108** (3), 565–615.

- EATON, J. and GROSSMAN, G. M. (1986). Optimal trade and industrial policy under oligopoly. *The Quarterly Journal of Economics*, **101** (2), 383–406.
- ERICSON, R. and PAKES, A. (1995). Markov-perfect industry dynamics: A framework for empirical work. *Review of Economic Studies*, **62**, 53–82.
- FABRA, N. and TORO, J. (2005). Price wars and collusion in the spanish electricity market. *International Journal of Industrial Organization*, **23** (3-4), 155–181.
- FOWLIE, M., REGUANT, M. and RYAN, S. (2016). Market-based emissions regulation and industry dynamics. *Journal of Political Economy*, **124** (1), 249–302.
- FRANKEL, J. A., VEGH, C. A. and VULETIN, G. (2014). On graduation from fiscal procyclicality. *Journal of Development Economics*, **100** (1), 32–47.
- GALI, J., GERTLER, M. and LOPEZ-SALIDO, J. (2007). Markups, gaps, and the welfare costs of business fluctuations. *Review of Economics and Statistics*, **89** (1), 44–59.
- GENESOVE, D. and MULLIN, W. P. (1998). Testing static oligopoly models: Conduct and cost in the sugar industry, 1890-1914. *RAND Journal of Economics*, **29** (2), 355–377.
- GIIGNL (2012). The lng industry in 2012, gIIGNL Annual Report.
- GOETZ, D. (2019). Broadband mergers and dynamic bargaining: An application to netflix. *Working Paper*.
- GOWRISANKARAN, G., NEVO, A. and TOWN, R. (2015). Mergers when prices are negotiated: Evidence from the hospital industry. *American Economic Review*, **105** (1), 172–203.

- GRENNAN, M. (2013). Price discrimination and bargaining: Empirical evidence from medical devices. *American Economic Review*, **103** (1), 145–177.
- GROSSMAN, M. G. (1990). Promoting new industrial activities: A survey of recent arguments and evidence, paris: OECD Economic Studies 14.
- GROSSMAN, S. J. and HART, O. D. (1986). The costs and benefits of ownership: A theory of vertical and lateral integration. *Journal of Political Economy*, **94** (4), 691–719.
- HALL, B. and VAN REENEN, J. (2000). How effective are fiscal incentives for r&d? a review of the evidence. *Research Policy*, **29** (4), 449–469.
- HANLON, W. (2018). The persistent effect of temporary input cost advantages in shipbuilding, 1850-1911.
- HANSEN, J. D., JENSE, C. and MADSEN, E. S. (2003). The establishment of the danish windmill industry. *Review of World Economics*, **139** (2), 324–347.
- HARDING, R. and SHEPPARD, D. (2017). Japan outlaws restrictions on resale of lng cargoes. *Financial Times*.
- HARRISON, A. and RODRIGUEZ-CLARE, A. (2010). Trade, foreign investment, and industrial policy for developing countries. *Handbook of Development Economics*, pp. 4039–4214.
- HARTLEY, P. (2015). The future of long-term lng contracts. *The Energy Journal*, **36** (3).
- HEAD, K. (1994). Infant industry protection in the steel rail industry. *Journal of International Economics*, **37**, 141–165.
- HENDEL, I. and NEVO, A. (2006). Measuring the implications of sales and consumer inventory behavior. *Econometrica*, **74**, 1637–73.

- HIRSCHMAN, A. O. (1958). *The Strategy of Economic Development*. Yale University Press.
- HO, K. and LEE, R. S. (2017). Insurer competition in health care markets. *Econometrica*, **85** (2), 379–417.
- HORN, H. and WOLINSKY, A. (1988). Bilateral monopolies and incentives for merger. *RAND Journal of Economics*, pp. 408–419.
- HORTACSU, A. and PULLER, S. L. (2008). Understanding strategic bidding in multi-unit auctions: a case study of the texas electricity spot market. *RAND Journal of Economics*, **39** (1), 86–114.
- HOTZ, J. and MILLER, R. (1993). Conditional choice probabilities and the estimation of dynamic models. *Review of Economic Studies*, **60**, 497–529.
- HUBBARD, T. N. (2000). The demand for monitoring technologies: the case of trucking. *Quarterly Journal of Economics*, **115** (2), 533–560.
- IEA (2013). *Developing a Natural Gas Trading Hub in Asia: Obstacles and Opportunities*. Paris: OECD/IEA.
- IGU (2009). 2006-2009 triennium work report, programme committee d study group 2, lng, chapter 3: Lng contracts.
- IRWIN, D. (2000). Could the united states iron industry have survived free trade after the civil war? *Explorations in Economic History*, **37**, 278–299.
- ITO, K. and REGUANT, M. (2016). Sequential markets, market power, and arbitrage. *American Economic Review*, **106** (7), 1921–1957.
- ITO, T. (1992). *The Japanese Economy*. MIT Press.

- ITSKHOKI, O. and MOLL, B. (2019). Optimal development policies with financial frictions. *Forthcoming in Econometrica*.
- JEON, J. (2018). Learning and investment under demand uncertainty in container shipping, department of Economics, The Boston University.
- JIN, H., QIAN, Y. and WEINGAST, B. R. (2005). Regional decentralization and fiscal incentives: Federalism, chinese style. *Journal of Public Economics*, **89**, 1719–1742.
- JOHNSON, C. (1982). *MITI and the Japanese Miracle: The Growth of Industrial Policy*. Stanford University Press.
- JOSKOW, P. (1985). Vertical integration and long-term contracts: The case of coal burning electric generation plants. *Journal of Law, Economics and Organization*, **1** (1), 33–80.
- (1987). Contract duration and relationship-specific investments: Empirical evidence from coal markets. *American Economic Review*, **77** (1), 168–185.
- KALOUPTSIDI, M. (2014). Time to build and fluctuations in bulk shipping. *American Economic Review*, **104** (2), 564–608.
- (2018). Detection and impact of industrial subsidies: The case of chinese shipbuilding. *Review of Economic Studies*, **85** (2), 1111–1158.
- , SCOTT, P. and SOUZA-RODRIGUES, E. (2018). Identification of counterfactuals in dynamic discrete choice models.
- KLEIN, B., CRAWFORD, R. G. and ALCHIAN, A. A. (1978). Vertical integration, appropriable rents, and the competitive contracting process. *The Journal of Law and Economics*, **21** (2), 297–326.

- KOYAMA, K. (2012). An analysis on asian premium for lng price, special Bulletin of The Institute of Energy Economics, Japan.
- KRUEGER, A. O. (1990). Government failures in development. *Journal of Economic Perspectives*, **4** (3), 9–23.
- KRUGMAN, P. (1991). *Geography and trade*. The MIT Press.
- KRUGMAN, P. R. (1986). *Strategic Trade Policy and the New International Economics*. Cambridge, MA: MIT Press.
- LANE, N. (2017). Manufacturing revolutions: Industrial policy and networks in south korea. *Working Paper of Institute for International Economic Studies, Stockholm*.
- LI, H. (2015). Intertemporal price discrimination with complementary products: E-books and e-readers, SSRN Working Paper.
- LIU, E. (2019). Industrial policies in production networks. *The Quarterly Journal of Economics*, **134** (4), 1883–1948.
- LIU, Y. and MAO, J. (2019). How do tax incentives affect investment and productivity? firm-level evidence from china. *American Economic Journal: Economic Policy*, **11** (3), 261–91.
- LUZIO, E. and GREENSTEIN, S. (1995). Measuring the performance of a protected infant industry: The case of brazilian microcomputers. *The Review of Economics and Statistics*, **77** (4), 622–633.
- MAGNOLFI, L. and RONCORONI, C. (2018). Political connections and market structure, department of Economics, The University of Wisconsin.

- MASTEN, S. E. and CROCKER, K. J. (1985). Efficient adaptation in long-term contracts: Take-or-pay provisions for natural gas. *American Economic Review*, **75** (5), 1083–1093.
- MILLER, N. and OSBORNE, M. (2014). Spatial differentiation and price discrimination in the cement industry: evidence from a structural model. *RAND Journal of Economics*, **45** (2), 221–247.
- and WEINBERG, M. C. (2017). Understanding the price effects of the millercoors joint venture. *Econometrica*, **85** (6), 1763–1791.
- NEUMANN, A. (2009). Linking natural gas markets—is lng doing its job? *The Energy Journal*, pp. 187–199.
- NEUMARK, D. and SIMPSON, H. (2015). Place-based policies. In G. Duranton, V. Henderson and W. Strange (eds.), *Handbook of Regional and Urban Economics*, vol. 5, Elsevier, Amsterdam.
- NEVO, A. (2001). Measuring market power in the ready-to-eat cereals industry. *Econometrica*, **69** (2), 307–342.
- and ROSSI, F. (2008). An approach for extending dynamic models to settings with multi-product firms. *Economics Letters*, **100**, 49–52.
- OECD (2015). Peer review of the korean shipbuilding industry and related government policies, paris: OECD.
- (2016). Peer review of the japanese shipbuilding industry, paris: OECD.
- OIES (2017). *The Outlook for Floating Storage and Regasification Units (FSRUs)*. Oxford Institute for Energy Studies.

- PAKES, A., OSTROVSKY, M. and BERRY, S. (2007). Simple estimators for the parameters of discrete dynamic games (with entry/exit examples). *RAND Journal of Economics*, **38** (2), 373–399.
- PONCET, S. (2005). A fragmented china: Measure and determinants of chinese domestic market disintegration. *Review of international Economics*, **13** (3), 409–430.
- POWELL, J. (1984). Least absolute deviations estimation for the censored regression model. *Journal of Econometrics*, **25** (1), 303–325.
- PULLER, S. L. (2007). Pricing and firm conduct in california’s electricity market. *The Review of Economics and Statistics*, **89** (1), 75–87.
- RITZ, R. (2014). Price discrimination and limits to arbitrage: An analysis of global lng markets. *Energy Economics*, **45** (1), 324–332.
- ROBINSON, P. M. (1982). On the asymptotic properties of estimators of models containing limited dependent variables. *Econometrica*, **50** (1), 27–41.
- RODRIK, D. (2010). The return of industrial policy. *Project Syndicate*.
- RUESTER, S. (2009). Changing contract structures in the international liquefied natural gas market: a first empirical analysis. *Revue d’économie industrielle*, (127), 89–112.
- RYAN, S. (2012). The costs of environmental regulation in a concentrated industry. *Econometrica*, **80** (3), 1019–1061.
- SALVO, A. (2010). Inferring market power under the threat of entry: the case of the brazilian cement industry. *RAND Journal of Economics*, **41** (2), 326–350.
- SEIM, K. (2006). An empirical model of firm entry with endogenous product-type choices. *Rand Journal of Economics*, **37:3**.

- STAPCZYNSKI, S. and MURTAGH, D. (2019). The future is now for lng as derivatives trading takes off. *Bloomberg*.
- STIGLITZ, J., LIN, J. Y. and MONGA, C. (2013). The rejuvenation of industrial policy, policy Research Working Paper 6628, The World Bank.
- STIGLITZ, J. E. and LIN, J. Y. (2013). *The Industrial Policy Revolution I: The Role of Government Beyond Ideology*. PALGRAVE MACMILLAN.
- STOPFORD, M. (2009). *Maritime Economics*. New York: Routledge.
- SWEETING, A. (2007). Market power in the england and wales wholesale electricity market 1995-2000. *Economic Journal*, **117** (1), 654–685.
- (2013). Dynamic product positioning in differentiated product markets: The effect of fees for musical performance rights on the commercial radio industry. *Econometrica*, **81** (5), 1763–1803.
- TALUS, K. (2011). Long-term natural gas contracts and antitrust law in the european union and the united states. *Journal of World Energy Law and Business*, **4** (3), 260–315.
- TERAZONO, E. (2019). Looser restrictions prompt surge in lng liquidity. *Financial Times*.
- THOMPSON, P. (2001). How much did the liberty shipbuilders learn? new evidence for an old case study. *Journal of Political Economy*, **109** (1), 103–137.
- (2007). How much did the liberty shipbuilders forget? *Management Science*, **53** (6), 908–918.
- TORNELL, A. and LANE, P. R. (1999). The voracity effect. *American Economic Review*, **89** (1), 22–44.

- UNCTAD (2018). *The Review of Maritime Transport 2018*. Geneva: UNCTAD secretariat.
- VILLAS-BOAS, S. B. (2009). An empirical investigation of the welfare effects of banning wholesale price discrimination. *RAND Journal of Economics*, **40** (1), 20–46.
- WILLIAMSON, O. E. (1975). Markets and hierarchies. *New York*, **2630**.
- (1979). Transaction-cost economics: The governance of contractual relations. *Journal of Law and Economics*, **22** (2), 233–61.
- WILSON, D. J. (2009). Beggar thy neighbor? the in-state, out-of-state, and aggregate effects of r&d tax credits. *The Review of Economics and Statistics*, **91** (1), 431–436.
- WOLAK, F. A. (2000). An empirical analysis of the impact of hedge contracts on bidding behavior in a competitive electricity market. *International Economic Journal*, **14** (2), 1–40.
- (2003). Lessons from the california electricity crisis, working Paper, Center for the Study of Energy Markets (CSEM).
- XU, D. (2008). A structural empirical model of r&d, firm heterogeneity, and industry evolution, working Paper.
- YI, F., LIN LAWELL, C.-Y. C. and THOME, K. E. (2015). A dynamic model of subsidies: Theory and application to the ethanol industry, working paper, Cornell University.
- YOUNG, A. (2000). The razor’s edge: Distortions and incremental reform in the people’s republic of china. *The Quarterly Journal of Economics*, **115** (4), 1091–1135.