From the holdings of the Department of Special Collections , Princeton University Library, Princeton, NJ

These documents are either in the public domain, or can only be used for educational and research purposes (Fair use) as per U.S. Copyright law (text below). By accessing this file, all users agree that their use falls in the public domain or within fair use as defined by the copyrightlaw. We do not charge any permission or usage fees for the publication of images of material inour collections, including those provided by Princeton University Library, via our website, catalog, or directly from staff.For more information regarding copyright, credit, and citation please review the information onthe Copyright, Credit and Citation Guidelines page of our website.
https://library.princeton.edu/special-collections/policies/copyright-credit-and-citation-guidelines. If you have any questions, including how to order higher-quality images from the collection, contact Special Collections Public Services via our Ask Us! Form. https://library.princeton.edu/special-collections/ask-us

---

Doctoral Dissertations, 2011-2026 Economics

2018

ESSAYS ON INTERMEDIATED MARKETS

Brancaccio, Giulia

Princeton, NJ : Princeton University

---

http://arks.princeton.edu/ark:/88435/dsp01p2676z258

U.S. Copyright law: The copyright law of the United States (Title 17, §108, United States Code) governs the making of photocopies or other reproductions of copyrighted material. Under certain conditions specified in the law, libraries and archives are authorized to furnish a photocopy or other reproduction. One of these specified conditions is that the photocopy or reproduction is not to be used for any purpose other than private study, scholarship, or research. If a user makes a request for, or later uses, a photocopy or reproduction for purposes in excess of fair use that user may be liable for copyright infringement. This institution reserves the right to refuse to accept a copying order if, in its judgment, fulfillment of the order would involve a violation of copyright law. https://www.copyright.gov/title17/title17.pdf

ESSAYS ON INTERMEDIATED MARKETS

GIULIA BRANCACCIO

A DISSERTATION

PRESENTED TO THE FACULTY

OF PRINCETON UNIVERSITY

IN CANDIDACY FOR THE DEGREE

OF DOCTOR OF PHILOSOPHY

RECOMMENDED FOR ACCEPTANCE

BY THE DEPARTMENT OF ECONOMICS

ADVISORS: PROFESSORS BO HONORÉ & JAKUB KASTL

JUNE 2018

© Copyright by Giulia Brancaccio, 2018.

All rights reserved.

Abstract:

In three essays, this dissertation examines the role of intermediaries in determining market outputs and trends in a variety of different settings.

Chapter 1 studies the motives for trade for financial intermediaries in OTC markets. In these markets, public information about fundamentals is limited, and trade takes place under conditions of asymmetric information. Financial intermediaries, then, may rely on their trading activity to acquire information about the state of market fundamentals. Information acquisition, therefore, becomes an additional motive for trade. In this paper, my co-authors and I use a combination of reduced-form techniques and structural analysis to characterize and measure experimentation motives for trade in the U.S. secondary market for municipal bonds. We find that financial intermediaries are willing to pay up to 15% of the intermediation spread to improve their information. Furthermore, we argue that experimentation explains up to 10% of the volume of trade in the market.

In Chapter 2, my co-authors and I leverage detailed data on vessel movements and shipping contracts to shed new light on how the world shipping industry affects world trade costs and trade flows. We build a framework for modeling how decisions of exporters and ships jointly determine trade costs and trade flows. This framework provides a novel link to understand trade patterns and we showcase this by studying the impact on trade flows of different trade shocks.

Finally, in Chapter 3, my co-authors and I investigate the importance of fuel oil costs in shaping world trade, through the shipping industry. We use a rich dataset on ships' movements and contracts, along with a dynamic model describing the world shipping industry, to measure the elasticity of trade with respect to ship fuel costs. We find that the average estimated elasticity is 0.35, but ranges from 0.1 to about 1.2 depending on the level of the fuel cost. Strikingly, the trade elasticity features a pronounced asymmetry in low vs. high oil prices, which is attributed to

the equilibrium of the transportation sector. Finally, we use the estimated elasticity to assess the importance of ship design on trade flows.

Acknowledgments:

This dissertation is the product of many advisers, faculty, students, and friends as much as it is of my own work. Foremost among them are my advisors Bo Honoré, Myrto Kalouptsidi, and Jakub Kastl. They have helped me through countless iterations of my research, and their encouragement and insights have been invaluable during my journey to become a researcher. Bo has mentored me since my first years at Princeton. He has always steered my work with grace and deep understanding. I hope I have learned that from him. Myrto's advice guided me to be the researcher that I didn't think I could be. She taught me to be pointed, thorough, confident, and to enjoy the ride. Jakub's input was crucial for my journey. His continued push to better myself has taught me to always strive for excellence.

Outside of my dissertation committee, I was also supported by the economics community at Princeton. In particular, I thank Nick Buchholz, Sylvain Chassang, Jan De Loecker, Adam Kapor, Eduardo Morales, and Motohiro Yogo for helpful comments and for challenging me at various stages of this dissertation. To my co-author Theodore Papageorgiou I would like to express my gratitude; this dissertation would have been much the lesser if not for his insights. Finally, this work benefited greatly from the stimulating environment provided by the graduate IO/Applied Micro group at Princeton University.

Without friends to share it with the Ph.D. would have been an impossible journey, and I thank Julia Fonseca, Clara Santamaria, and Molly Schnell for keeping me sane throughout these crazy years.

I owe a deep and lifelong debt of thanks to my mother, for instilling in me a love of learning and the teaching me the value of hard work; to my sister, for always being there for me; to Mauro, for never giving up.

Finally, I thank Tommaso for always helping me to be the best version of myself.

To Tommaso, of course

Contents

Abstract iii

Acknowledgments v

1 Learning by Trading: The Case of the U.S. Market for Municipal Bonds 1

1.1 Introduction 1

1.2 Industry and Data Description 9

1.2.1 The Secondary Market for Municipal Bonds 9

1.2.2 Data 11

1.3 Reduced-Form Evidence 13

1.3.1 Learning Through Inter-Dealer Trade 13

1.3.2 Improvement in Market Transparency 16

1.4 Model 21

1.4.1 Environment 21

1.4.2 Behavior 28

1.5 Empirical Strategy 36

1.5.1 Experience 36

1.5.2 Dealers' Information 40

1.5.3 Dealers' Costs 42

1.6 Results 44

1.6.1 First stage 45

1.6.2 Dealers' Costs . . . 47
1.7 Value and Precision of Information . . . 48
1.7.1 Precision of Information . . . 49
1.7.2 Value of Information . . . 52
1.8 Implications for Market Transparency . . . 56
1.9 Conclusion . . . 60
References . . . 60
1.A Additional Tables and Figures . . . 65
1.B What do Dealers Know? . . . 66
1.C Definition of Θ . . . 67
1.D Assets Classes . . . 69
1.E Choice Probabilities . . . 71
1.F Estimation Algorithm . . . 74
1.G Steady-State conditions . . . 76
1.H Proofs from Section 1.5 . . . 78

2 Geography, Search Frictions and Endogenous Trade Costs 81
2.1 Introduction . . . 81
2.2 Industry and Data Description . . . 89
2.2.1 Trade in Dry Bulk Commodities . . . 89
2.2.2 Data . . . 90
2.3 Facts . . . 92
2.3.1 Trade Flows, Geography and Ballasting . . . 93
2.3.2 Trade Costs (Shipping Prices) . . . 95
2.3.3 Search Frictions . . . 96
2.4 Model . . . 102
2.4.1 Environment . . . 102
2.4.2 Equilibrium . . . 106

2.5 Empirical Strategy . . . 113
2.5.1 Matching Function Estimation . . . 113
2.5.2 Ship Costs . . . 118
2.5.3 Exporter Valuations and Costs . . . 120
2.6 Results . . . 122
2.6.1 Matching Function . . . 122
2.6.2 Ship Costs . . . 125
2.6.3 Exporter Valuations and Costs . . . 127
2.7 Counterfactuals . . . 131
2.7.1 Change in Trade due to a Change in Shipping Efficiency . . . 131
2.7.2 Chinese Slow-down . . . 133
2.7.3 Opening the Northwest Passage . . . 135
2.7.4 The Trade Lost because of Search Frictions . . . 137
2.8 Conclusion . . . 138
References . . . 138
2.A Construction of Ship Travel Histories . . . 144
2.B Construction of Searching Ships . . . 146
2.C Additional Figures and Tables . . . 146
2.D Proof of Proposition 2.4.2 . . . 151
2.E Proof of Proposition 2.5.1 . . . 152
2.F Proof of Lemma 2.5.2 . . . 154
2.G Identification of Ship Port and Sailing Costs . . . 155
2.H Algorithm for computing the steady state equilibrium . . . 157

3 The Impact of Oil Prices on World Trade 159

3.1 Introduction . . . 159
3.2 Industry and Data . . . 164
3.3 Model . . . 166

3.3.1 Environment 166
3.3.2 Equilibrium 167
3.3.3 Model Estimation and Results 170
3.4 Fuel Cost and Trade Elasticity 176
3.5 Oil Shocks and Price Dispersion 182
3.6 Fleet Fuel Efficiency and World Trade 184
3.7 Conclusions 188
References 188

Chapter 1

Learning by Trading: The Case of the U.S. Market for Municipal Bonds[1]

1.1 Introduction

The desire to acquire information provides a powerful explanation for many economic phenomena. Often information is explicitly purchased by, for example, investing in market research or hiring experts. In contexts where agents learn from the consequences of their actions, the demand for information creates a trade-off between choosing the best action for today and choosing the most informative action about tomorrow: to acquire valuable information, agents might decide to deviate from the action myopically most profitable and "experiment."

This trade-off is particularly relevant in markets where trade is decentralized. Here, the lack of a centralized trading mechanism often implies that public information about prices and volume of trade is limited. Therefore, agents must rely on their own trading activity to acquire information about the market fundamentals: negotiating with others reveals information about the counterparty's valuation of the

[1]The paper on which this chapter is based is co-authored with Dan Li and Norman Schürhoff

asset. This, in turn, provides valuable information about the overall state of the market.

In this paper, we investigate whether information acquisition is a first-order motive for trade in decentralized markets and explore the implications of experimentation for the functioning of these markets. Our setting is the secondary market for U.S. municipal bonds, a decentralized market with trades totaling over $3 trillion per year. Two reduced-form facts reveal that information acquisition is a key driver of trading decisions for agents operating in this market. These facts motivate us to build a framework for studying information acquisition as a motive for trade. Forward-looking dealers build an inventory of an asset by trading with one another (inter-dealer trade) and with myopic investors. The investors' valuation for the asset depends on a persistent and unobserved common preference shock. Trading is costly but reveals information about investors' valuations, creating incentives to experiment. To quantify these incentives to experiment, we estimate the model using rich micro-level data on trading activity. The results suggest that information is valuable: we find that dealers are willing to pay up to 15% of the intermediation spread to double the precision of their information about market fundamentals. Moreover, we study the process of information diffusion in the market and find that dealers acquire information mostly by trading among themselves, rather than directly trading with investors. Finally, we find that experimentation explains up to 10% of the volume of trade in the market.

States and municipalities throughout the United States depend on the municipal bond market to raise funds for investments in schools, highways, and other public projects. Several features of this market make it an ideal laboratory to study the interaction between trading and experimentation. First, trade in municipal bonds is decentralized, and learning about the terms of trade involves participating in trade directly. Second, a large number and variety of bonds are outstanding at any given

time, and each asset includes a variety of complex special provisions. This complicates pricing and makes information acquisition a first-order issue. Finally, in recent years, the Municipal Security Rulemaking Board (MSRB) —the regulatory body for the municipal bond market— has taken a number of concrete steps to improve access to information about market activity for market participants. These changes in the information structure that agents face allow us to directly test the importance of experimentation in explaining the behavior of financial institutions active in the market.

We use a proprietary data set from the MSRB that covers the universe of transactions involving a municipal bond between 2000 and 2005. Importantly, the data contain an identifier for the dealers involved in each transaction, allowing us to construct the complete trading history for each dealer.

Two reduced-form facts illustrate the relevance of information acquisition for shaping the trading behavior of agents operating in decentralized markets. First, data on inter-dealer trade show that after a dealer sells an asset to another dealer at a particularly high price, he will increase the price he charges to his clients for the same asset. A variety of placebo tests suggest that this result is not spurious and indicates that this change in behavior is likely driven by information acquisition. Trade, therefore, can a be source of valuable information.

Second, we look at the outcome of a policy change that increased access to public information about trading activity and prices in the municipal bond market. We focus on uninsured assets where incomplete information is arguably more severe, since insurance protects investors against the risk of default. We find that trade between dealer and investors for assets that are uninsured falls compared to insured assets. Moreover the price at which dealers are willing to buy (sell) uninsured assets decreases (increases). This suggests that information acquisition is a key incentive that prompts dealers to take on the risk of holding this type of asset in inventory.

To rationalize these facts, we build a dynamic model of trading in decentralized markets where dealers trade and experiment. Forward-looking dealers build costly inventories of municipal bond holdings by trading with myopic retail investors and other dealers. Investors' valuations for the asset change over time due to a persistent, common, and unobservable preference shock (the "market fundamental"). Since dealers can choose the timing of trade, their returns depend critically on the information they are able to acquire about the state of the market. Incentives to experiment enter the decision to trade: trading is costly but trading prices are informative about the unobserved shock. When facing retail investors, each trade has the same information content, but trading with more investors allows the dealer to sample more observations and it is both riskier and costlier. When trading with other dealers, trading prices are informative about the counterparty's valuation for the asset, which in turn reveals what he knows about the state of the market. Some dealers have better information than others due to their trading history. To capture the decentralized nature of trade, we allow dealers to only observe a summary statistic of the past trading activity of their peers, which we call "experience."

A dealer's trading decisions depend on two unobserved objects: their information about fundamentals and their experience. In the estimation, we first use inter-dealer trading and trading history to recover the dealer's information about the state of fundamentals and experience. Next, we use the dealer's trading choices to recover the core set of primitives —trading costs and the cost of holding inventory.

We exploit an implication of the model to recover dealers' experience: more experienced dealers will pay lower prices to buy assets in the inter-dealer market. Our baseline specification compares how prices for trades executed by a specific seller, in a specific month and asset, change depending on the past trading experience of the buyer. We focus on comparisons for a fixed month and seller to ensure that the estimates are robust with respect to market-wide shocks. We also consider

alternative specifications to address the potential bias introduced by unobserved buyer's heterogeneity. In particular, we use inflow and outflow of funds in the market for municipal bonds to build dealer-specific liquidity shocks and include these as instruments for the dealer's experience.

To recover the dealers' information about the fundamentals, we assume that dealers only acquire information through trade or through public signals accessible to everyone, the econometrician included. We exploit a Hansen-Sargan test for over-identifying restrictions to show that dealers have no information about the state of fundamentals of an asset in periods in which they did not trade the asset. This suggests that learning activities in the market for municipal bonds are strongly connected to "realized" trade and provides a justification for this assumption.

The estimated model serves two purposes. First, we use the model to characterize the value and precision of information for dealers active in the market. We find that dealers are willing to pay up to 15% of their average intermediation spread (i.e., the difference between the price at which they buy and sell the asset) to double the precision of their estimate of the asset's market value. Furthermore, we find that experimentation allows dealers to increase the precision of this estimate by 25%. Finally we study how information is disseminated across agents.

Second, we explore the impact of experimentation on the volume of trade in the market. We find that improving market transparency can reduce the volume of trade for an asset by more than 10% by weakening the incentives to experiment. Two effects are at play. On the one hand, transparency weakens the incentives to experiment. On the other hand, it reduces uncertainty about the state of the market fundamental. This makes dealers more confident and gives them incentives to trade larger quantities of the asset, partially offsetting the first effect. The final balance of these two forces varies dramatically depending on the assets' underlying primitives. For this reason

we perform a comparative static analysis to identify the features of the asset that will determine the success of this types of policies.

We focus on volume of trade for a number of reasons. Many authors, starting with Pagano 1989, Kyle 1985, 1989, and Admati and Pfleiderer 1988, have argued that volume of trade is a key variable to determine whether investors can sell an asset on short notice without loss (that is, to determine whether an asset is "liquid"). Issuers, in turn, pay a high price to issue illiquid securities. As an example, Wang et al. 2008 estimate that municipal bond issuers pay $13 billion a year to compensate investors for the risks implied by the illiquidity of the market. Increasing the volume of trade in this market, therefore, can translate to huge savings for local governments and municipalities. Finally, volume of trade can be an interesting outcome variable per se, as historically it has been the target of policies addressing the inefficiencies of decentralized financial markets.

Related Literature This paper is at the intersection of three principal strands of literature. The basic trade-off between learning and sacrificing immediate payoff is focal in the literature on strategic experimentation. We empirically quantify the strength and implications of this basic trade-off in the context of decentralized financial markets. Finally, we integrate these concepts with ideas from empirical studies of industry dynamics.

Experimentation has long been studied in economics, mostly from a theoretical standpoint (for a survey, see Hörner and Skrzypacz 2016). Several papers within this literature explicitly share our focus on experimentation as a motive for trade —most notably Aghion et al. 1993, Grossman et al. 1977, Mirman et al. 1993, and Kihlstrom et al. 1984. Our focus remains an empirical one. For this reason, we strip the incentives to experiment to their minimal components. This makes the agents' problem tractable, allowing us to bring the model to the data.

Several papers, such as Leach and Madhavan 1993 and Bloomfield and O'Hara 1999, 2000, have discussed the implications of experimentation for the trading behavior of agents in financial markets. Furthermore, Wolinsky 1990, as well as Golosov et al. 2014, and Blouin and Serrano 2001 explore the linkages between trading and information diffusion in a decentralized market with private information. Their objective is to study under what conditions all relevant information is revealed over time. Despite this interest, direct quantification of the role of experimentation and measurement of its implication for market structure has remained scarce. We contribute to this literature by employing a tractable analytical framework to empirically study the role of incentives to experiment as a motive for trade.

We integrate the literature on experimentation with a recent literature that uses search models to study the trading behavior of agents in decentralized markets. Largely, these papers build on the framework developed by Duffie et al. 2005, 2007, to study search frictions in the context of these markets both in theory (most recently, Hugonnier et al. 2014 and Farboodi et al. 2016) and empirically (Gavazza 2011b,a). In this paper, we focus on a different feature of decentralized markets: the lack of public information about trade activity.[2] For this reason, we borrow the most basic structure of these models and enrich it with incomplete information and learning; in our setup the decision to trade not only depends on inventory management and search costs but also on experimentation.

Finally, we relate to the literature on industry dynamics (e.g., Hopenhayn 1992, Ericson and Pakes 1995) which characterizes Markov-perfect equilibria in entry, exit, and investment choices given some uncertainty in the evolution of the states of firms and their competitors. Instead of these choices, we model the agents' problem as a series of trading and pricing decisions. Since agents interact with one another repeatedly, this problem generates a particularly high-dimensional state space. We

[2]Duffie and Manso 2007 have a similar focus, but they focus on information diffusion rather than information acquisition.

introduce a number of innovations to mitigate the computational burden. First, to simplify trading decisions, we assume that agents only observe a summary statistic of the trading history of other market participants. This assumption not only permits solving for equilibrium policies of agents and simplifies the agents' inference, but also reflects a more realistic behavioral model for decentralized markets. Moreover, given the large number of dealers in the market, we assume that the distribution of dealers' private states is perfectly forecastable by agents, conditional on the preference shocks they are trying to learn. This approach has precedent in the literature on firm dynamics (Weintraub et al. 2008). Finally, our empirical methodology borrows from the literature on the estimation of dynamic setups (e.g., Rust 1987, Aguirregabiria and Mira 2007, Bajari et al. 2007, and Pakes et al. 2007) in exploiting conditional choice probabilities to obtain information on the value functions and, in turn, on the primitive of interest.

The rest of the paper is structured as follows. Section 1.2 provides a description of the industry and the data used. Section 1.3 presents a number of reduced-form facts suggesting that experimentation is a first-order determinant of agent's trading decisions in the market. Section 1.4 describes the model. Section 1.5 lays out our empirical strategy, while in Section 1.6 we present the estimation results and characterize dealers' incentives to experiment. Section 1.8 discusses the implication of our results in term of market structure, while Section 1.9 concludes. The Appendix contains additional tables and figures, proofs to our propositions, as well as further data and estimation details.

1.2 Industry and Data Description

1.2.1 The Secondary Market for Municipal Bonds

Municipal bonds are debt securities issued by states, cities, and other local governments to fund day-to-day obligations and to finance capital projects. Their importance cannot be overstated: in 2017 they were the main source of funding for 75% of the total public investment in infrastructure.

To ease credit access for local governments, interest rates accrued on municipal bonds are exempt from individual income taxes both at the federal and the local level. Due to this obvious tax advantage, 70% of the total municipal debt outstanding is held by private investors directly (50.2%) or through mutual funds (20%).

The secondary market for municipal bonds is organized as a standard decentralized market. There is no central exchange for municipal securities, and financial institutions registered with the SEC as municipal securities broker-dealers intermediate trades among investors. Dealers execute nearly all transactions in a "principal capacity": the dealers buy the assets directly and hold them in inventory until they are able to find a buyer. Dealers, moreover, can trade among themselves in the inter-dealer market. Every year there are more than 2,000 active broker-dealers, and the largest market share is around 10%.

At any given time there is a large number of bonds outstanding, each of which includes complex features. This lack of standardization worsens incomplete information and makes information acquisition a first-order issue. In particular, over our sample period there are 1.5 million different assets outstanding, issued by more than $50,000$ different units of state and local governments.[3] Moreover, several types of special provisions can complicate pricing. As an example, callable bonds are redeemable by the issuer before the scheduled maturity, while a sinking fund provision requires

[3]For comparison, this is 20 times the number of corporate bond types.

the issuer to retire a specified portion of debt each year. Furthermore, nonstandard interest payment frequencies are not uncommon, and most of the outstanding assets have some form of credit enhancement.[4] The majority of the outstanding assets have more than one of these special provisions: Harris and Piwowar 2006 show that only a small fraction of the outstanding assets (around 14%) contain no complexity features.

The lack of a centralized trade mechanism together with the lack of standardization imply that the information needed to price the assets is often not public. Yet, some coarse indexes about market activity are publicly available. In particular, since 1995, the MSRB has published information about the volume of trade and average trading price for assets traded more than four times during the previous day ("next-day reporting"). This, however, covers only 5% of the assets traded. Moreover, the most widely watched municipal bonds indexes are compiled by "*The Bond Buyer*." These indexes are either based on dealers' estimates for the price of a hypothetical bond[5] or on the activity on the primary market. However, these are too coarse to effectively reduce the uncertainty for the pricing of individual bonds.

Access to public information about trade activity has improved steadily in recent years. The four-trade threshold for next-day reporting was abandoned on June 23, 2003, when all trades began to be reported the next day. Moreover, since January 31, 2005, information about each transaction in the market is made available online within 15 minutes of the execution of a trade. Investors seem to have embraced the new source of information with enthusiasm: on the first day of 15-minute trade reporting, The Bond Market Association reported that the website on which trades were reported averaged about 10,000 visits per minute.[6] In Section 1.3 we leverage these changes to show that incentives to experiment are a first-order determinant of the volume of trade in the market.

[4] An issuer improves the credit rating of a security by purchasing the financial guarantee (e.g., insurance, letter of credit) of a large financial intermediary.

[5] This is the case for the 20-Bond Index, the 11-Bond Index, and the Revenue Bond Index.

[6] See Schultz 2012.

Finally, investors' participation in the secondary market is driven by liquidity shocks rather than speculation. Municipal bonds are considered to be a relatively safe investment, with historically low default rates. As an example, for Aa- and A-rated municipal bonds, the 10-year cumulative default rate is .03% compared to 0.8% for corporate bonds. The low default risk and the composition of the owners, tend to make municipal bonds "buy-and-hold" investments. In other words, municipal bonds are mainly bought at issuance and held until maturity. For this reason, when we think of incentives to experiment we focus on dealers' incentives to trade. Consistent with this, trades on the secondary market are small: the median trade is worth $\$25,000$, and 80% of trades have a value of less than $\$100,000$.

1.2.2 Data

Our main data source is the proprietary Transaction Reporting System audit trail from the MSRB. In an effort to improve market transparency, the MSRB has required dealers to report all transactions in municipal securities since 1998. The transactions data cover the 5-year period from January 2000 to December 2005. For every transaction involving municipal bonds, our data provide information about the terms of the trade, such as the trading price, date and time of the trade as well as par value (the value at maturity of the asset exchanged, or the volume of the trade) of the asset, and an asset identifier. Significantly, we observe identifiers for the dealer firm intermediating each trade: for customer trades, the data identify the dealer buying and the dealer selling the bond, while for trades among dealers, the data identify the dealers on each side of the trade. In addition to the comprehensive transactions data, we obtained reference information on all municipal bonds, including issuance date, maturity, coupon, taxable status, ratings, call features, issue size, and issuer characteristics from Thomson SDC. Finally, we obtain the time series for market bond indexes, as well as monthly municipal mutual fund flows from Bloomberg.

We filter the transactions to eliminate data errors and ensure data completeness. For a bond to be in our sample, it must have complete descriptive data in the SDC and satisfy a number of trade-specific filters and bond specific filters (fixed or zero coupon, non-derivative, non-warrant, not puttable, maturity $\geq$1 year, \$5K denomination). Since the focus of this paper is on the secondary market, we remove all trades during the first 90 days after issuance and less than one year away from maturity.[7]

Summary statistics Our final data set involves $20,207,244$ trades on the secondary market between 2000 and 2005, involving $587,224$ unique assets. As shown in Table 2.1, on average 34 million dollars worth of assets are bought or purchased by private investors every month. The average price is \$99.45, across sales and purchases, with substantial variation (the overall standard deviation is \$10.68, and the median standard deviation within each month is \$10.44). The difference between the price paid to and from investors within a month (the "intermediation spread") is on average 2%. Consistent with the description of the market in Section 1.2.1, the trade size is on average \$70,000 (the median is \$25,000) and institutional size trades (above 1 million) happen sporadically (they represent 1% of the total trades).

There are $4,072$ different dealers active in the market over our sample period. The largest dealer intermediates 10% of total trades, while the second largest dealer has less than a 5% market share. We obtain a similar picture if we use a narrower definition of "market", that takes into account the possibility that dealers specialize. For instance, the highest market share by state of issuance is on average 15%.

Interaction on the inter-dealer market is sparse: the inter-dealer trade is one-third that of trade with investors. Finally, dealers' trade relationships do not seem to be very persistent. As can be seen in the last row of Table 2.1, on average (across dealers) every second transaction on the inter-dealer market involves a new counterparty.

[7] As a result of these filters, we retain 65% of all the transactions included in the initial dataset.

Table 1.1: Summary statistics

	Mean	St. Dev	Median	Min	Max
Trading Price	99.484	10.68	101.52	36.644	116
Intermediation spread	2.1	1.55	1.19	-0.23	6.8
Monthly trade to investors (10^7 USD)	3.45	0.51	3.38	2.35	4.85
Monthly inter-dealer trade (10^7 USD)	1.50	0.25	1.48	1.02	2.25
Trade size ($1,000$ USD)	72.05	190.92	25	5	2,245
Dealers' market share	0.043	0.40	0.00026	2^{-7}	11.6
Inter-dealer trades with a new counterparty (%)	44.49	37.63	30.66	0.07	1

Notes: The above table provides summary statistics for trading activity on the secondary market for muncipal bonds US. Data come from the proprietary Transaction Reporting System audit trail from the MSRB,and covers the universe of transaction in this market between 2000 and 2005.

1.3 Reduced-Form Evidence

In this section, we present reduced-form evidence that suggests that: (i) dealers acquire information through trade; and (ii) dynamic incentives to acquire information are an important determinant of dealers trading and pricing behavior.

1.3.1 Learning Through Inter-Dealer Trade

In markets where public information about trading activity is limited, agents need to rely on private interactions with other agents to aggregate the information dispersed in the market. In particular, negotiating with others can reveal information about the counterparty's valuation for the asset. This, in turn, provides valuable information about the overall state of the market.

We use data on inter-dealer trade to argue that dealers do extract information from prices in inter-dealer trades and change their trading behavior to account for

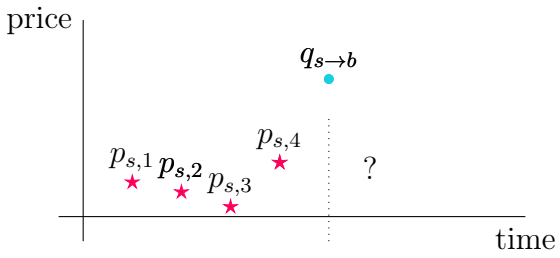

Figure 1.1: Experiment

this information. This suggests that bargaining, and trade in general, can be a source of valuable information for market participants.

We consider the situation depicted in Figure 1.1. In particular, consider a dealer s who sells an asset to a dealer b at price $q_{s\to b}$. Suppose that price $q_{s\to b}$ is higher than the average price that dealer s was charging to his clients in the previous week, $(p_{s,1}, \ldots, p_{s,4})$. Seller s could interpret price $q_{s\to b}$ as a signal that the asset is more valuable than he thought. In this case, one might expect that he will revise his pricing strategy and increase the price he charges to his clients.

Concretely, let i denote a generic inter-dealer trade, for asset a_i at price q_i[8]. Moreover, s_i and b_i, denote, respectively, seller and buyer involved in trade i. Similarly, let j denote a generic *sale* to an investor, executed by dealer s_j, at price p_j. For every inter-dealer trade i we construct the average price charged by seller s_i, for asset a_i, to his clients in the two weeks preceding (or following) the trade, $\hat{p}_i^{\text{before}}$ (or $\hat{p}_i^{\text{after}}$), as

$$\hat{p}_i^{\text{before}} = \frac{\sum_{j\geq 1} p_j \mathbb{I}\left\{-2 \leq \text{week}_j - \text{week}_i < 0, a_i = a_j, s_i = s_j\right\}}{\sum_{j\geq 1} \mathbb{I}\left\{-2 \leq \text{week}_j - \text{week}_i < 0, a_i = a_j, s_i = s_j\right\}}.$$

We summarize the change in dealer s_i pricing behavior after trade i with the quantity $\Delta_i = \frac{\hat{p}_i^{\text{after}} - \hat{p}_i^{\text{before}}}{\hat{p}_i^{\text{before}}}$. The orange density in Figure 1.2 plots differences Δ_i for all those trades i for which $\hat{p}_i^{\text{before}} < q_i$. Remarkably, after 87% of such trades, dealer s_i increases the price he is charging to his (non-dealer) clients. On average, prices change by 1%, and this average change is significant at the 5% level. The second density in Figure

[8]Here and throughout the paper, q denotes the price on the inter-dealer market.

1.2 shows the differences Δ_i for all those trades i for which $\hat{p}_i^{\text{before}} > q_i$. Remarkably, in this case, only 55% of the prices increase, and the change is not significant.

Figure 1.2: Change in pricing behavior after inter-dealer trade

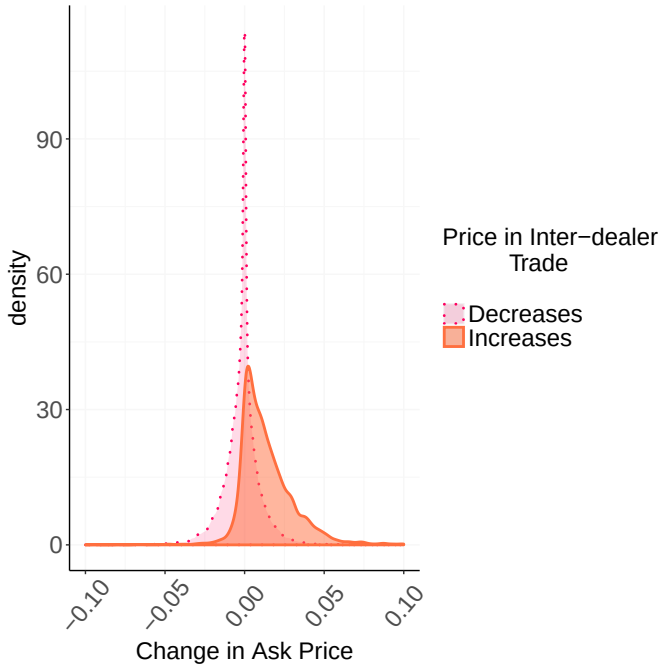

Notes: the above figure plots the histogram of the difference of the average price charged by dealers to retail investors before and after inter-dealer trades. The orange density considers trades in which average price $\hat{p}_i^{\text{before}}$ is lower than inter-dealer trading price q_i. Instead, the second density consider trades of the type described in Fig 1.1.

Figure 1.3 plots the results of two different placebo tests to verify that this result is not spurious. First, one might worry that the result in 1.2 captures liquidity shocks that lead sellers s_i to increase prices charged both to other dealers and to private investors. In this case, one would expect that the dealer will also increase the prices charged for other assets, not necessarily traded on the inter-dealer market. For this reason, we check whether the sellers s_i involved in the trades in Figure 1.2 also change the prices they are charging for assets that they have not traded in the inter-dealer market. In particular, for every inter-dealer trade i considered in Figure 1.2 we

construct

$$p^{\text{before}}_{\text{placebo liq},i} = \frac{\sum_{j\geq 1} p_j \mathbb{I}\left\{-2 \leq \text{week}_j - \text{week}_i < 0, a_i \neq a_j, s_i = s_j\right\}}{\sum_{j\geq 1} \mathbb{I}\left\{-2 \leq \text{week}_j - \text{week}_i < 0, a_i \neq a_j, s_i = s_j\right\}}.$$

The right panel of Figure 1.3 plots the difference $\Delta_{\text{placebo liq},i} = \frac{p^{\text{after}}_{\text{placebo liq},i} - p^{\text{before}}_{\text{placebo liq},i}}{p^{\text{before}}_{\text{placebo liq},i}}$ for all the trades i included in Figure 1.2. In this case, only 50% of the prices increase and the average change is not significantly different from zero.

In the same fashion, one might worry that the change in behavior captured by Figure 1.2 is driven by a market-wide shock to the value of asset a_i. If this were true, one would expect that also dealers not participating in inter-dealer trades will change the price they are charging for asset a_i. To verify whether this is the case, we construct

$$p^{\text{before}}_{\text{placebo mkt },i} = \frac{\sum_{j\geq 1} p_j \mathbb{I}\left\{-2 \leq \text{week}_j - \text{week}_i < 0, a_i = a_j, s_i \neq s_j\right\}}{\sum_{j\geq 1} \mathbb{I}\left\{-2 \leq \text{week}_j - \text{week}_i < 0, a_i = a_j, s_i \neq s_j\right\}}.$$

The left panel of Figure 1.3 plots the difference $\Delta_{\text{placebo mkt},i} = \frac{p^{\text{before}}_{\text{placebo mkt},i} - p^{\text{before}}_{\text{placebo mkt},i}}{p^{\text{before}}_{\text{placebo mkt},i}}$ for all the trades i included in Figure 1.2. Once again, only 48% of the prices increase, and the average change is not significantly different from zero.

1.3.2 Improvement in Market Transparency

For years the SEC has been warning private investors and Congress about the need to improve access to information about trade activity in the market for municipal bonds. This pressure from the SEC culminated in a series of provisions aimed at improving market transparency. In particular, on June 23, 2003, the MSRB started distributing daily summaries about the trading activity in the market during the previous day. Moreover, starting on January 31, 2005, the MSRB mandated that

Figure 1.3: Placebo tests

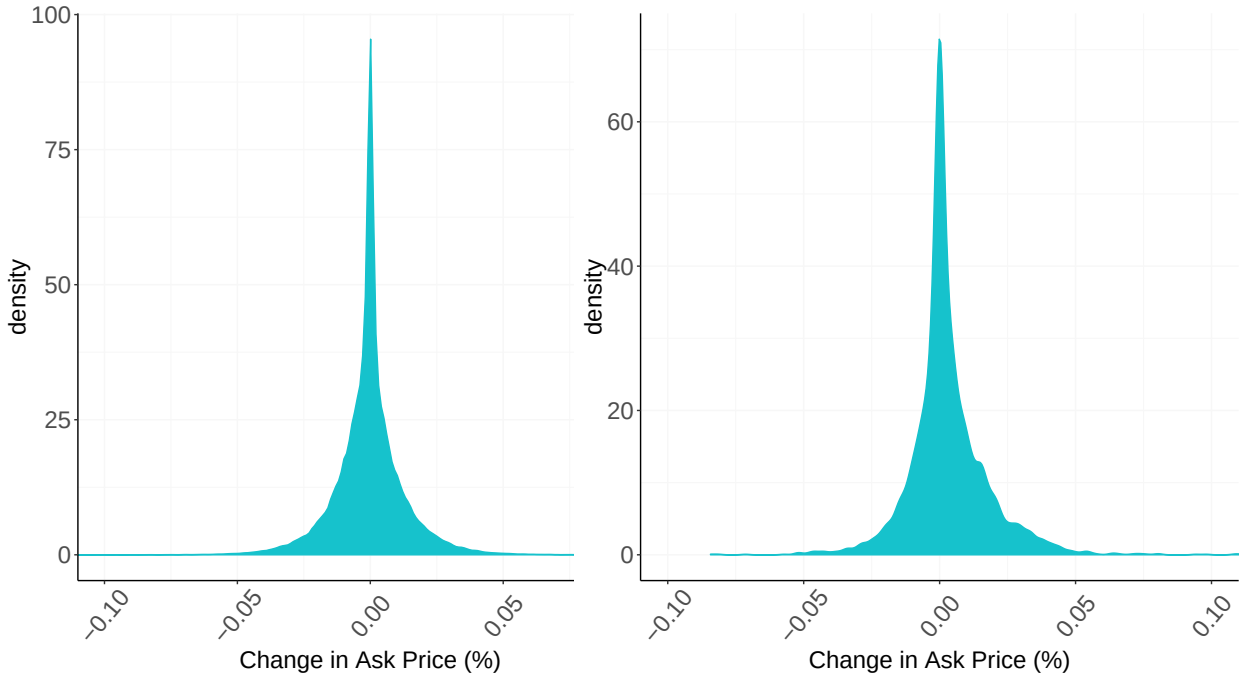

Notes: the above figure plots the results of two different placebo tests. To test whether the result is driven by dealer-specific liquidity shocks, in the right panel we plot changes in the seller's pricing behavior for assets not traded on the inter-dealer market. To test whether the result is driven by asset-specific demand shocks, in the left panel we plot changes in pricing behavior for dealers not participating to inter-dealer trades.

details of all transactions in U.S. municipal bonds be reported on a timely basis and posted online almost immediately.[9]

Proponents of market transparency argue that the lack of public information about trading activity gives dealers an informational advantage. Dealers, the argument goes, exploit this advantage to extract rents from their clients by "selling high and buying low". Market transparency, by leveling the playing field, would increase investors' participation and benefit the market at large. For instance, SEC commissioner Arthur Levitt remarked, "The undeniable truth is that transparency helps investors make

[9] Asquith et al. 2013 study the effect of a similar policy intervention in the market for corporate bonds, and find similar results.

better decisions, and it increases confidence in the fairness of the markets. And, that means more efficient markets, more trading, more market liquidity."[10]

This argument, however, ignores *dealers'* incentives to trade. Information acquisition motives for trade, in particular, can substantially erode the positive effects of transparency. Indeed, when public information about market activity is limited, trading with investors allows dealers to acquire valuable information about the market value of the asset. This generates an additional motive for trade that market transparency might weaken. Therefore, if information acquisition is a key determinant of trading and pricing decisions for financial intermediaries, improving access to public information might result in a decrease of trading activity, as well as a worsening of trading prices for investors.

We explore the effect of the 2003 policy change through a difference-in-difference set-up. We leverage the idea that improving transparency will have stronger consequences for assets for which incomplete information is more severe. A typical example of these assets in the market for municipal bonds are uninsured assets. Issuers that meet certain credit criteria can purchase municipal bond insurance policies from large private insurance companies. The insurance guarantees the payment of principal and interest on a bond issue if the issuer defaults. Pricing for insured assets, therefore, is more straightforward compared to pricing for the uninsured ones and depends less on unobserved factors.

We focus on two main outcome variables. First, we look at the response of trading activity, which we measure as the number of bonds traded times the par value exchanged in each trade (i.e. the value at maturity of the asset exchanged). Next, we look at the impact of information dissemination on the trading conditions for investors. In particular, we focus on the difference between the average ask and bid price within a week (the "intermediation spread"). From an investor's standpoint,

[10]Speech before the Bond Market Association.

the intermediation spread represents the out-of-pocket transaction costs of trading an asset. Instead, from a dealer's standpoint, it affects the incentive to participate in trade.

We estimate Equation

$$y_{it} = \gamma_{0,i} + \gamma_1 t + \gamma_2 \mathbb{I}\{t > t_0\} + \gamma_3 x_i t + \lambda x_i \mathbb{I}\{t > t_0\} + \epsilon_{it}, \tag{1.1}$$

where y_{it} is issue i's outcome on week t, t_0 is the week in which the policy change is implemented, $\gamma_{0,i}$ is an indicator for bond i, and x_i is an indicator for whether the asset is uninsured. In Equation 1.1, any pre-existing difference between assets is captured by $\gamma_{0,i}$, while the effects of the policy that accrue to all bonds are absorbed by coefficient γ_2. The coefficient of interest is λ, which estimates the effect of transparency on trading outcomes for uninsured assets. Finally parameter γ_3 absorbs any potential pre-existing trend for uninsured assets.

Table 1.2 reports estimates of the parameters in Equation 1.1 for different outcomes and time windows.

First, we focus on the effect of the policy on volume of trade between dealers and investors. The estimate of the effect of the introduction of transparency on uninsured assets is negative and significant. The volume of trade for uninsured assets drops by 2.8% compared to insured assets in the 6-month window around dissemination, which is significant at the 1% level.

Next, we turn to the intermediation spread for trades between dealers and retail investors. The estimates in the second half of Table 1.2 show that the intermediation spread for uninsured assets increases compared to insured assets. This pattern reinforces the conclusion that the decrease in volume of trade depends on the weakening

of information acquisition motives for trade and suggests that a key incentive for dealers to trade these assets is information acquisition.[11]

Table 1.2: Difference-in-difference estimates

	Volume of Trade (log)			Intermediation Spread		
	6 Months	9 Months	12 Months	6 Months	9 Months	12 Months
uninsured * $\mathbb{I}\{t > t_0\}$	-0.028^{***}	-0.025^{***}	-0.021^{***}	0.020^{***}	0.020^{**}	0.020^{***}
	(0.010)	(0.008)	(0.008)	(0.007)	(0.008)	(0.007)
N	1,438,297	2,142,769	2,847,241	228,741	335,994	451,482
Level	Issuer-Week					

$^{***}p \leq 0.01$, $^{**}p \leq 0.05$, $^{*}p \leq 0.1$

Notes: The above table presents the output from the difference-in-difference regression that measures the effect of the change in market transparency on trading activity (first three columns) and intermediation spreads (last three columns). We use insured assets as control group. Observations are at the asset-week level, and standard errors are clustered at the asset level.

It is worth emphasizing that the estimates presented in Table 1.2 provide reduced form evidence that experimentation motives to trade are key. However, reduced form approaches cannot control for changes in the behavior of investors and, therefore, are not sufficient to directly measure the impact of market transparency on dealer's incentives to trade. In Section 1.8 we leverage our model to isolate the effect of market transparency on volume of trade through its impact on incentives to experiment. This not only allows us to quantify the role of experimentation but also to identify the possibly opposing channels through which incentives to experiment affect dealers' incentives to trade. Looking at how the model primitives shape this different channels, makes it possible to identify the critical features that determine the success of the policy for a specific class of assets.

[11] In Appendix 1.A we show that mean trade size doesn't change for uninsured assets as a result of the increase in transparency. This suggests that adverse selection is not the key determinant of the change in intermediation spread and volume of trade.

1.4 Model

In this section, we introduce a tractable dynamic model of trade in decentralized markets. Our goal is to capture dealers' incentive to experiment. We begin by describing agents' characteristics and objectives, as well as the interaction between trade and experimentation. In Section 1.4.2 we study players' dynamic choice problem.

1.4.1 Environment

Time $t \in \{1, 2, \ldots\}$ is discrete and a unique asset is traded. The market is populated by two types of agents: short-run investors and long-run dealers. Everyone is risk neutral.

Investors are myopic and live only for one period. Each investor is either a buyer or a seller. His valuation for the asset depends on an idiosyncratic component and a common preference shock $\theta_t \in \Theta$. Common shock θ_t represents common factors that affect investors' willingness to pay for the asset, and it is unobserved by investors and dealers alike. Each investor knows his own valuation but doesn't understand its correlation with other investors' valuations.

Common shock θ_t evolves over time according to a discrete Markov chain; denote by $h(\theta_t|\theta_{t-1})$ the probability of moving from state θ_{t-1} to θ_t. Publicly available information about the common shock is summarized by public signal y_t^P observed in each period t with mean θ_{t-1}. In the context of the market municipal bonds, this public signal captures information contained in monthly indexes about the market performance, as well as information about the performance of municipal mutual funds.

Dealers $d \in \{1, \ldots, D\}$ are forward-looking players with time preferences determined by a constant discount rate $\beta > 0$. In every period, dealers can trade the asset with investors and among themselves. Assets bought and not sold accumulate over time and form inventory $x_{d,t} \in \{0, 1, \ldots\}$. Due to the illiquidity of the market for

municipal bonds, short selling is rare and costly. For this reason, we assume that short selling is infeasible. Therefore inventory is always positive. Carrying inventory is costly: in every period dealers pay a cost $\kappa(x_{d,t}) \geq 0$. Inventory cost $\kappa(\cdot)$ captures frictions that prevent dealers to increase balance sheet size, such as the cost of capital (usually inventory is levered) or limits to exposure to risk. The only source of revenue for the dealers is the resale price of the asset, while operating costs depend on the price paid to buy the assets and costs required to carry out trades.

Over time dealers form beliefs about the common shock. We denote by $\pi_{d,t} \in \Delta(\Theta)$ the probability that dealer d assigns to different values of θ_t given the information he has accumulated before time t.

We begin describing the trading protocol for trade with investors. Then we describe then we detail the trading protocol for inter-dealer trade.

Trade with Investors. In every period t, each dealer d decides whether to buy or sell (or do nothing) the asset to investors and how many investors to search for and trade with. We focus on the dealers' role in inter-temporal intermediation, rather than cross-sectional intermediation. For this reason we assume that dealers cannot both buy and sell the asset to investors in each period. This assumption is natural in the market of municipal bonds, where for only 7% of dealer-month-issuer pairs we observe a "comparable"[12] amount of sales and purchases. In each trade only one unit of the asset is exchanged.

Denote by $n_{d,t} \in \{-x_{d,t}, \ldots, N\}$ the number of units of the asset the dealer trades with investors. The lower bound $n_{d,t} \geq -x_{d,t}$ comes from the assumption of no short sales. The upper bound is for notational convenience.[13]

[12] That is, for 7% of triplets composed by a dealer d, month t and issuer a, it is $\frac{\max\{\text{sales}_{a,d,t}, \text{purhases}_{a,d,t}\}}{\text{sales}_{a,d,t}+\text{purhases}_{a,d,t}} < 0.75$.

[13] Since cost of trade is convex, this assumption is without loss of generality.

Trading with investors is costly. The cost of trading captures separately factors that affect the dealer's decision on whether to buy or sell the asset, and on how many units of the asset to trade. Let $\left(\epsilon_{d,t}^{\text{buy}}, \epsilon_{d,t}^{\varnothing}, \epsilon_{d,t}^{\text{sell}}\right) \in \mathbb{R}^3$ be a cost shock i.i.d. over time and across dealers. For trading $n_{d,t}$ units of the asset dealer d pays a fixed cost $c_F\left(n_{d,t}, \epsilon_{d,t}\right)$ depending on whether he buys or sells the asset: given cost parameters $c^{\text{buy}}, c^{\text{sell}} \in \mathbb{R}$,

$$c_F\left(n_{d,t}, \epsilon_{d,t}^{\text{buy}}, \epsilon_{d,t}^{\varnothing}, \epsilon_{d,t}^{\text{sell}}\right) = \left(c^{\text{buy}} + \epsilon_{d,t}^{\text{buy}}\right) \mathbb{I}_{n_{d,t}>0} + \epsilon_{d,t}^{\varnothing} \mathbb{I}_{n_{d,t}=0} + \left(c^{\text{sell}} + \epsilon_{d,t}^{\text{sell}}\right) \mathbb{I}_{n_{d,t}<0}.$$

Parameters c^{buy} and c^{sell} can be negative to capture fees that the dealer might demand from his clients. Furthermore, let $\epsilon_{d,t}^{1} \in \mathbb{R}$ be a cost shock i.i.d. over time and across dealers, the dealer also pays a search cost $c_V\left(n_{d,t}, \epsilon_{d,t}\right)$ to find investors interested in trading: given $c_1, c_2 \geq 0$,

$$c_V\left(n_{d,t}, \epsilon_{d,t}^{1}\right) = c_1 \left|n_{d,t}\right| + c_2(n_{d,t} \cdot n_{d,t}) + \epsilon_{d,t}^{1} \left|n_{d,t}\right|.$$

Letting $\epsilon_{d,t} = \left(\epsilon_{d,t}^{\text{buy}}, \epsilon_{d,t}^{\varnothing}, \epsilon_{d,t}^{\text{sell}}, \epsilon_{d,t}^{1}\right)$, the overall trading cost is

$$c\left(n_{d,t}, \epsilon_{d,t}\right) = c_F\left(n_{d,t}, \epsilon_{d,t}^{\text{buy}}, \epsilon_{d,t}^{\varnothing}, \epsilon_{d,t}^{\text{sell}}\right) + c_V\left(n_{d,t}, \epsilon_{d,t}^{1}\right).$$

.

The price $p_{i,d,t} \geq 0$ received or paid by dealer d in each trade with investors $i \in \{1, \ldots, |n_{d,t}|\}$ is the outcome of a bargaining process between dealers and investors. We abstract away from the specifics of this bargaining process and capture its outcome in a reduced form way: trading prices are i.i.d. draws from a distribution $f\left(\cdot|\theta_t, \text{sign}\left(n_{d,t}\right)\right)$ which depends on the current realization of the unobserved state, θ_t, and on whether the dealer is buying or selling, $\text{sign}\left(n_{d,t}\right)$.[14] This modeling approach

[14] $\text{sign}\left(n_{d,t}\right)$ equals 1 if $n_{d,t} > 0$, it equals -1 if $n_{d,t} < 0$, and it is 0 otherwise.

allows us to specify the relevant variable for experimentation —prices' informational content— in a parsimonious way.[15]

In sum, given cost shock $\epsilon_{d,t}$ and trading prices $(p_{i,d,t})_{i=1}^{n_{d,t}}$, the payoff from trading $n_{d,t}$ units of the asset is

$$c\left(n_{d,t}, \epsilon_{d,t}\right) - \text{sign}\left(n_{d,t}\right) \sum_{i=1}^{|n_{d,t}|} p_{i,d,t}.$$

Experimentation enters the decision of how many units of the asset to trade, since trading prices are noisy signals about θ_t. In particular, observing the trading prices allows the dealer to acquire information about the state of the market. This information is valuable, since it allows the dealer to anticipate changes in the resale value for the asset and, therefore, to improve the future timing of his trading decisions. In Section 1.4.2 we describe more in detail how dealers update after trading with investors.

Trade with dealers. After trading with investors, dealers can trade with one another. Inter-dealer trade proceeds as follows. A constant share α of the dealers is randomly selected to be "potential sellers."[16] Each potential seller d can make a take-it-or-leave-it offer, $q_{d,t} \geq 0$, to trade one unit of the asset. In contrast with the random search literature, the offer from potential sellers is directed to a specific dealer $\tilde{d}$ on the other side of the market, a "potential buyer." We allow dealers to exchange only one unit of the asset in the context of inter-dealer trade. Indeed, we capture

[15] An alternative way to think about this modeling choice is that, consistent with Green et al. 2010, we are assuming that dealers have strong market power vis-a'-vis retail investors. Under these circumstances dealers are able to extract all the surplus when trading with an investor and the trading price coincides with the investor's valuation.

[16] The role of α is similar to that of the number of potential entrants, as it defines an upper bound on the total volume of trade in the market. The total quantity traded in the inter-dealer market remains endogenous since dealers can decide not to engage in trade, and buyers can reject the offer received.

the dealers' decision on the extensive margin of trade in the context of trade with investors, since inter-dealer trade represents one-third of total volume trade.

Potential buyers can either accept one of the offers received or reject them all. This assumption is rarely binding for the market for municipal bonds since, conditional on participating to inter-dealer, for 80% of dealer-month pairs we see the dealer buying assets from only one counterparty. Making an offer is costly: $c^{\text{d2d}}\left(\tilde{d}, \xi_{d,t}\right)$ is the cost of making an offer to dealer $\tilde{d}$ given i.i.d. cost shock $\xi_{d,t}$. This shock accounts for idiosyncratic reasons that might lead a dealer to favor one counterparty over another. For example, the dealer might find the line occupied. Accepting the offer costs nothing.[17]

After a pair of dealers d and $\tilde{d}$ have traded, they exchange information on what they know about θ_t. We model this as follows: Dealer d observes a (possibly noisy) signal $y_{\tilde{d},t}$ of dealer $\tilde{d}$'s current belief $\hat{\pi}_{\tilde{d},t} \in \Delta(\Theta)$ about θ_t. Belief $\hat{\pi}_{\tilde{d},t}$ is simply $\pi_{\tilde{d},t}$ updated after the trades with investors.[18] We don't model the dealers' strategic decision about the information to reveal. This approximation is reasonable since the communication happens *after* trade and since we are working under the assumption that the market is large, and therefore the probability that dealers will interact again in the future is small.[19]

In the inter-dealer market, each dealer d is characterized by his level of *experience* $e_{d,t} \in \{1, \ldots, E\}$, a publicly observed summary statistic of his history of trades up to time t. Experience is a proxy for the precision of each dealer's information about common shock θ_t. Dealers accumulate experience by trading with retail investors and

[17] We cannot estimate separately the cost born by the buyer and the seller, since the resulting probability of trade depends on jointly on both. However, we also estimate the model assuming that the buyer makes the offer (and pays the cost) to the seller. The results are similar.

[18] After trade with investors, dealers update their beliefs according to Bayesian updating. See Section 1.4.2 for the details.

[19] Duffie and Manso 2007 adopts a similar model for how information is exchanged.

more experienced counterparties in the inter-dealer market. Denote by

$$r\left(e_{d,t}|n_{d,t-1}, e_{\tilde{d},t-1}, e_{d,t-1}\right) \tag{1.2}$$

the probability that dealer d has experience $e_{d,t}$ given that in the previous period he had experience $e_{d,t-1}$, then traded with $n_{d,t-1}$ investors and with dealer $\tilde{d}$ of experience $e_{\tilde{d},t-1}$. Given past history, experiences are drawn independently across dealers. We also assume that all levels of experience are recurrent; this is consistent with the idea that experience may depreciate over time. Finally, r is increasing (with respect to first order stochastic dominance) in its arguments; this captures the idea that dealers with a richer trading history have more precise information about common shock θ_t.

Experience is the main determinant of trading decisions in the inter-dealer market. Dealers are ex-ante homogeneous and differences in valuation for the asset and in information about θ_t emerge over time only because of differences in trading history. For this reason, potential sellers choose whom to trade with on the basis of what they know about the past trading activity of their peers —dealers who have been trading more will have more precise information about common shock θ_t. However, assuming that the entire past history of trades is commonly known would be not only computationally cumbersome, but also unrealistic given the opacity of the market for municipal bonds. A public summary static like experience is a parsimonious solution to these issues. In Section 1.5 we describe how we define dealers' experience in the data.

Beyond experience, inter-dealer trade is "*anonymous*": dealers do not keep track of the identity of their trading counterparties. This assumption is natural since in the market for municipal bonds interaction on the inter-dealer market is sparse.[20] Consistently, we assume that the cost to make an offer only depends on the experience

[20]In the market for municipal bonds, the median number of interaction between two dealers, conditional on them interacting at all, is 3.

of the recipient, and not on his identity. If dealer $\tilde{d}$ and $\tilde{\tilde{d}}$ have the same level of experience, then

$$c^{\text{d2d}}\left(\tilde{d}, \xi_{d,t}\right) = c^{\text{d2d}}\left(\tilde{\tilde{d}}, \xi_{d,t}\right).$$

More specifically, potential seller d observes i.i.d. cost shocks $\xi_{d,t} = \left(\xi^{\tilde{e}}_{d,t}, \xi^{0}_{d,t}\right) \in \mathbb{R}^{\bar{E}+1}$ and pays cost $c^{\text{d2d}}(\tilde{e}) + \xi^{\tilde{e}}_{d,t}$ to make an offer to a dealer with experience $\tilde{e}$, and $\xi^{0}_{d,t}$ if he decides not to trade.

By trading with one another, dealers acquire information about one another's information about θ_t: this is another way to experiment. Both the offer received by potential buyers and the reply received by potential sellers, as well as post-trade communication, convey information about what the counterparty knows about θ_t. To avoid the infinite regress problem of learning what others know what others know..., we make the following simplifying assumption. Borrowing from the literature on social learning,[21] we assume that each dealer behaves as if the information received from any other dealer $\tilde{d}$ is *independent* of what he already knew, conditional on the realization of state θ_t and the dealer's experience, $e_{\tilde{d},t}$. This is a reasonable assumption in the context of a large market where dealers share a common history of trades with very low probability.

Finally, we assume that potential buyers and sellers only update their beliefs based on *realized trade*. This assumption is driven mainly by empirical concerns. Our data does not show offers to sell the asset that were rejected by the buyer, and therefore we cannot identify changes in dealers' beliefs that derive from offers to trade that were rejected. This assumption is consistent with anecdotal evidence which suggests

[21] It is standard in the social learning literature to assume that agents learn through DeGroot rules of thumb models, which often involve double-counting information. Most notably, Ellison and Fudenberg 1993, 1995 are benchmarks for the rule of thumb learning models. Moreover Chandrasekhar et al. 2012 exploit an experimental setup to argue that a DeGroot rule of thumb model of learning might provide a better description of agents learning on a network, compared to standard bayesian updating.

that there are strong reputational concerns involved in soliciting quotes only for their informational content, without the actual intention to buy or sell the asset. Moreover, in Appendix 1.B we use an Hansen-Sargan test for over-identifying restrictions to show the results of a test suggesting that learning activities in the market for municipal bonds are strongly connected to "realized" trade. This suggests that the empirical bite of this assumption is limited.

Timing. To summarize, in each period the timing is as follows:

1. θ_t is realized. Dealers observe public signal y_t^P about last period's shock θ_{t-1}. Then, they update their beliefs both to account for y_t^P and to account for the evolution of θ from the last period, according to h. Finally each dealer d pays a cost which depends on accumulated inventory $x_{d,t}$;

2. Dealers can trade with investors. Each dealer draws i.i.d cost shocks $\epsilon_{d,t} = \left(\epsilon_{d,t}^{\text{buy}}, \epsilon_{d,t}^{\text{sell}}, \epsilon_{d,t}^{\varnothing}, \epsilon_{d,t}^{1}\right)$ and decides with how many investors $n_{d,t} \in \{-x_{d,t}, \ldots, N\}$, if any, to trade with. Trading prices are i.i.d draws from a distribution $f\left(\cdot|\theta_t, \text{sign}\left(n_{d,t}\right)\right)$, which depends on the current value of the common preference shock θ_t. Dealers interpret prices as noisy signals about θ_t.

3. Dealers can trade with one another. The population is randomly divided among potential buyers and sellers. Potential sellers can make a t.i.o.l.i offer to a potential buyer to buy one unit of the asset, and potential buyers can either accept one of the offers received or reject all of them. After trade, dealers exchange information on what they know about θ_t.

1.4.2 Behavior

We first spell out the updating rules that dealers use to incorporate information obtained in the context of trade. Next, we derive the optimal behavior of dealers,

as well as equilibrium prices. Where it does not generate confusion, we drop the d subscript and use "tilde" to denote state variables of dealer d's trading counterparty. For instance we use $\tilde{e}$ instead of $e_{\tilde{d}}$ to denote the experience of dealer d's trading counterparty, $\tilde{d}$.

In this paper, we focus on a steady state of the model such that: (i) the fraction of dealers with a given inventory, belief, and experience depend on θ but not time; and (ii) the fraction of dealers with a given experience is constant in θ and time. This assumption is natural in the market at hand, where more than 2,000 dealers are active and each of them intermediates less than 10% of the total trade. For this reason, below we drop the dependence of value functions and choice probabilities on the vector e_t.

Updating Dealers first acquire information about common shock θ_t from prices in trades with investors. Dealer d, after observing prices $\vec{p}_{n,t} = (p_{i,t})_{i=1}^{|n|}$ updates according to standard Bayesian updating:

$$\hat{\pi}_{d,t}\left(\theta_t = \theta^k | \vec{p}_{n,t}\right) = \frac{f\left(\vec{p}_{n,t} | \theta^k, \text{sign}\,(n)\right) \pi_{d,t}\left(\theta^k\right)}{\sum_\theta f\left(\vec{p}_{n,t} | \theta, \text{sign}\,(n)\right) \pi_{d,t}\left(\theta\right)} := \mathcal{L}_{inv}\left(\pi_{d,t}; \vec{p}_{n,t}\right). \tag{1.3}$$

Dealers also acquire information about common shock θ_t by interacting on the inter-dealer market. Each pair of dealers d and $\tilde{d}$ involved in a trade, first observe the action of their counterparty, and then observe signal y. The inference dealers draw from these depends on their conjecture about the counterparty's private history, conditional on common shock θ_t. The assumption of *independence* implies that dealer d's conjecture does not depend on his own private state. The assumption of *anonymity* implies that this conjecture does not depend on the identity of the counterparty.

Denote $f^*\left(\tilde{y}, \tilde{q}, \tilde{e} | e, \theta\right)$ dealers' conjecture about the probability that a dealer with experience level $\tilde{e}$ will make offer $\tilde{q}$ and communicate signal $\tilde{y}$ to a dealer with experience level e, conditional on common shock θ. Symmetrically, $f^*\left(y, \text{accept} | \tilde{e}, q, e, \theta\right)$

denotes dealers conjecture about the probability that a dealer with experience $\tilde{e}$ will accept offer q and communicate signal y to a dealer with experience e, conditional on state θ. Seller d, with beliefs $\hat{\pi}_{d,t}$, after trading with buyer $\tilde{d}$ updates according to

$$\begin{aligned}\hat{\hat{\pi}}_{d,t}\left(\theta_t=\theta^k|y_{\tilde{d},t},e_{\tilde{d},t},q_{d,t},e_{d,t}\right)&=\frac{f^*\left(y_{\tilde{d},t},\text{accept}|e_{\tilde{d},t},q_{d,t},e_{d,t},\theta^k\right)\hat{\pi}_{d,t}\left(\theta^k\right)}{\sum_\theta f^*\left(y_{\tilde{d},t},\text{accept}|e_{\tilde{d},t},q_{d,t},e_{d,t},\theta\right)\hat{\pi}_{d,t}\left(\theta\right)}\\&:=\mathcal{L}_{\text{sell}}\left(\pi_{d,t};y_{\tilde{d},t},e_{\tilde{d},t},q_{d,t},e_{d,t}\right).\end{aligned}\tag{1.4}$$

Buyer $\tilde{d}$, instead, updates according to

$$\begin{aligned}\hat{\hat{\pi}}_{\tilde{d},t}\left(\theta_t=\theta^k|y_{d,t},q_{d,t},e_{d,t},e_{\tilde{d},t}\right)&=\frac{f^*\left(y_{d,t},q_{d,t},e_{d,t}|e_{\tilde{d},t},\theta^k\right)\hat{\pi}_{\tilde{d},t}\left(\theta^k\right)}{\sum_\theta f^*\left(y_{d,t},q_{d,t},e_{d,t}|e_{\tilde{d},t},\theta\right)\hat{\pi}_{\tilde{d},t}\left(\theta\right)}\\&:=\mathcal{L}_{\text{buy}}\left(\hat{\pi}_{\tilde{d},t};y_{d,t},q_{d,t},e_{d,t},e_{\tilde{d},t}\right).\end{aligned}\tag{1.5}$$

Note that since transition matrix r is increasing in past trading activity, more experienced dealers have more information about common shock θ_t. Therefore, signal $y_{d,t}$ communicated by more experienced dealers will be (possibly weakly) more informative than that communicated by dealers with lower experience level. This is consistent with the interpretation of experience as a proxy for the dealers' precision of information about common shock θ_t.

Trade with investors Let $V(\pi,x,e,\epsilon)$ be the value at the beginning of the period for a dealer who observes cost shock ϵ and has private history (π,x,e) . Then V satisfies

$$\begin{aligned}&V(\pi,x,e,\epsilon)\\&=-\kappa(x)+\max_{n\in\{-x,\ldots,N\}}\left\{-c(n,\epsilon)-\text{sign}(n)\,\mathbb{E}\left(\sum_{i=1}^{|n|}p_{it}|\pi,\text{sign}(n)\right)+\mathbb{E}\left[W\left(\mathcal{L}_{inv}\left(\pi,\vec{p}_n\right),x'(n;x),e,n\right)\right]\right\}.\end{aligned}\tag{1.6}$$

At the beginning of the period, the dealer pays a cost that depends on the inventory owned $\kappa(x) \geq 0$, and decides how many investors n to search for. Trading prices vis-a'-vis investors depend on the type of trade (sale or purchase), as well as on the unobserved asset's value θ. Not knowing θ, the dealer forms expectations

$$\mathbb{E}\left(\sum_{i=1}^{|n|} p_{it}|\pi, \text{sign}(n)\right) = \sum_{\theta\in\Theta} \pi(\theta) \sum_{i=1}^{|n|} \int p_i f(p_i|\theta, \text{sign}(n))\, dp_i.$$

The dealer's private state changes after trading. In particular, the dealer's beliefs about the current value of θ evolves according to updating rule 1.3, while his inventory evolves according to $x'(n;x) = x + n$. Finally, the dealer moves on to inter-dealer trade, where he obtains value W, which is defined below.

Next, we characterize the dealer's policy function following Kalouptsidi 2014. For tractability, we work under the assumption that : (i) cost shocks $\left(\epsilon_{d,t}^{\mathbf{buy}}, \epsilon_{d,t}^{\mathbf{sell}}, \epsilon_{d,t}^{\varnothing}\right) \in \mathbb{R}^3$, are drawn from a double exponential distribution F_0, with standard deviation σ_0; (ii) $\epsilon_{d,t}^1 \in \mathbb{R}$ is drawn from a normal distribution F_1 with standard deviation σ_1; (iii) $\tilde{c}(n) = c_1|n| + c_2 n^2$ is convex in $|n|$; and (iv) experience transition matrix r can be rewritten as

$$r(e'|e, n, \tilde{e}) = \sum_{e'} r_{\text{d2d}}(e'|e'', \tilde{e})\, r_{\text{inv}}(e''|e, n),$$

where r_{d2d} and r_{inv} describe, respectively, the change in experience that can be attributed to inter-dealer trade and to trade with investors. Under these conditions the dealer first decides whether to buy, sell, or avoid trading. Then he decides how many units of the asset to trade, comparing each trading level n to $n+1$ and $n-1$. Denote by $V^{\text{sign}(n)}(\pi, x, e, \epsilon^1)$ the dealer's highest utility conditional on either buying

or selling the asset:

$$V^{\text{sign}(n)}\left(\pi, x, e, \epsilon^1\right)$$
$$= \max_{n \in \mathcal{N}(\text{sign}(n))} \left\{ -\tilde{c}\left(n\right) - \epsilon_1 \left|n\right| - \mathbb{E}\left(\sum_{i=1}^{|n|} p_{it} | \pi, \text{sign}\left(n\right) \right) + \mathbb{E}\left[W\left(\mathcal{L}_{inv}\left(\pi, \vec{p}_{n,t}\right), x'\left(n; x\right), r_{\text{inv}}\left(n, e\right) \right) \right] \right\},$$

where

$$\mathcal{N}\left(\text{sign}\left(n\right)\right) = \begin{cases} \{1, \ldots, N\} & \text{sign}\left(n\right) = +1 \\ \{-x, \ldots, -1\} & \text{sign}\left(n\right) = -1 \end{cases},$$

The probability that dealer d chooses to trade with $n \neq 0$ investors can be written as

$$P\left(n_{d,t} = n | \pi, x, e\right) = \int_{\text{lb}(\pi,x,e,n)}^{\text{ub}(\pi,x,e,n)} \frac{\exp\left(\frac{V^{\text{sign}(n)}\left(\pi,x,e,\epsilon^1\right) - W(\pi,x,e)}{\sigma_0} \right)}{\exp\left(\frac{V^{\text{sign}(n)}(\pi,x,e,\epsilon^1) - W(\pi,x,e)}{\sigma_0} \right) + \exp\left(\frac{V^{\text{sign}(n)}(\pi,x,e,\epsilon^1) - W(\pi,x,e)}{\sigma_0} \right) + 1} dF_1\left(\epsilon^1\right), \tag{1.7}$$

where $\text{ub}\left(\pi, x, e, n\right)$ and $\text{lb}\left(\pi, x, e, n\right)$ are optimal policy thresholds defined in Appendix 1.E.

With a standard abuse of notation below we denote

$$V\left(\pi, x, e\right) = \mathbb{E}\left[V\left(\pi, x, e, \epsilon\right)\right].$$

Inter-dealer trade. Consider the situation of a potential buyer with private history (π, x) and experience e, who receives an offer to buy a unit of the asset at price $\tilde{q}$, from a dealer with experience $\tilde{e}$. The dealer decides whether to accept the offer by comparing the value from purchasing the asset at price $\tilde{q}$, and from rejecting

the offer:

$$\tilde{W}^{\text{buy}}(\pi, x, e; \tilde{q}, \tilde{e}) \tag{1.8}$$
$$= \max\left\{-\tilde{q} + \beta\mathbb{E}\left[V\left(\mathcal{L}_{\text{buy}}(\pi; \tilde{y}, \tilde{q}, \tilde{e}, e), x', r_{\text{d2d}}(\tilde{e}, e)\right) | e, \tilde{e}, \tilde{q}\right], \beta\mathbb{E}\left[V\left(\pi, x, r_{\text{d2d}}(0, e)\right)\right]\right\}$$

If the dealer accepts the offer, he pays price $\tilde{q}$, his inventory evolves according to $x' = x + 1$, his beliefs evolve following transition 1.5, and his experience evolves according to transition 1.2.

Updating rule $\mathcal{L}_{\text{buy}}(\pi; \tilde{y}, \tilde{q}, \tilde{e}, e)$ depends on the signal $\tilde{y}$ that the seller will communicate after trading. When deciding whether to accept offer $\tilde{q}$, the seller computes the expectation of this signal, conditional on offer $\tilde{q}$ and experience $\tilde{e}$ of his counterparty. Significantly, the offer will depend both on the counterparty's inventory as well as on his beliefs. For this reason, the distribution of signal $\tilde{y}$ conditional on offer $\tilde{q}$ is non-degenerate. If the dealer decides to reject the offer, his experience depreciates and he moves on to the next period, where he will obtain value $\beta\mathbb{E}\left[V\left(\pi, x, r_{\text{d2d}}(0, e)\right)\right]$ from trading with investors.

Next, consider the situation of a potential seller with private type (π, x) and experience e, who decides to make an offer to a potential buyer with experience $\tilde{e}$. The offer $q(\pi, x, e, \tilde{e})$ solves

$$\tilde{W}^{sell}(\pi, x, e; \tilde{e}) = \max_{q} \mathbb{P}(accept\ q | e, \tilde{e})\, \beta\mathbb{E}\left[V\left(\mathcal{L}_{sell}(\pi; \tilde{y}, \tilde{e}, q, e), x', r_{d2d}(\tilde{e}, e) | e, \tilde{e}, q\right)\right] \tag{1.9}$$
$$+\mathbb{P}(reject\ q | e, \tilde{e})\, \beta\mathbb{E}\left[V\left(\pi, x, r_{d2d}(0, e)\right)\right].$$

If the offer is accepted the seller's beliefs evolve according to transition 1.4, his inventory changes to $x' = x - 1$, and his experience evolves according to transition 1.2.

Finally, consider the decision of dealer d about the identity of his trading counterparty. The assumption of anonymity allows us to rephrase the potential seller's decision as that of choosing the optimal level of experience of the buyer to whom the offer should be sent:

$$\begin{aligned} & W^{\text{sell}}(\pi, x, e, n; \xi) \\ & = \max\left\{\max_{\tilde{e}\in\{1,\ldots,E\}}\left\{-c^{\text{d2d}}(\tilde{e}) + \tilde{W}^{\text{sell}}(\pi, x, e; \tilde{e}) + \xi^{\tilde{e}}\right\}, \beta\mathbb{E}\left[V\left(\pi, x, r_{\text{d2d}}(0, e)\right)\right] + \xi^{0}\right\}. \end{aligned} \tag{1.10}$$

The potential seller decides whether to propose trade to a potential buyer or to move to the next period and obtain value $\beta\mathbb{E}\left[V\left(\pi, x, r_{\text{d2d}}(0, e)\right)\right]$. If instead he decides to make an offer to a player of type $\tilde{e}$, he pays cost $c^{\text{d2d}}(\tilde{e}) + \xi^{\tilde{e}}$, and obtains value $\tilde{W}^{\text{sell}}(\pi, x, e, \tilde{e})$, defined in 1.9. We assume that shocks $\xi \in \mathbb{R}^{E+1}$ are draws from a double exponential distribution F_{ξ}, with standard deviation σ_{ξ}. This implies that the probability that a dealer with state (π, x, e) makes an offer to a dealer with experience $\tilde{e}$ satisfies

$$P_e(\tilde{e}|\pi, x, e) = \frac{\exp\left(\frac{-c(\tilde{e})+\tilde{W}^{\text{sell}}(\pi,x,e,\tilde{e})}{\sigma_{\xi}}\right)}{\exp\left(\frac{\beta\mathbb{E}[V(\pi,x,r_{\text{d2d}}(0,e))]}{\sigma_{\xi}}\right) + \sum_{\tilde{e}}\exp\left(\frac{-c(\tilde{e})+\tilde{W}^{\text{sell}}(\pi,x,e,\tilde{e})}{\sigma_{\xi}}\right)}. \tag{1.11}$$

Equation 1.11 and 1.10 can be interpreted as an approximation to the discrete choice problem faced by the dealer. Indeed, as cost shock variance σ_{ξ} converges to zero, $P_e(\tilde{e}|\pi, x, e)$ approaches 1 for level of experience $\tilde{e}$ for which utility $-c(\tilde{e}) + \tilde{W}^{\text{sell}}(\pi, x, e, \tilde{e})$ is largest.

Equilibrium. Dealers' policy functions depend on their current private and public history (π, x, e), as well as on beliefs about the policies of competitors. Other dealers' beliefs and inventory are unobservable, consistent with the decentralized nature of trade and the opacity of the market. Dealers, therefore, do not observe the valuation for the asset and the policy functions of their peers, but rather have beliefs about

these. Beliefs over other dealers' policies and valuations determine dealers behavior in the context of inter-dealer trade. Similarly to Weintraub et al. 2008, we assume that dealers' conjectures about their peers' private state are anchored to their long-run distribution. To allow for learning in the context of inter-dealer trade, however, dealers' conjecture depend on the long run distribution of dealers' private state *conditional* on the unobserved common preference shock θ_t. An equilibrium for the market described in the previous section is a distribution $f^*(\pi_d, x_d, e_d|\theta)$ of beliefs π_d, inventory x_d, and experience e_d across dealers conditional on preference shock θ, such that:

1. The distribution of experience in the population does not depend on θ

$$\begin{aligned} f_e^*(\bar{e}|\theta) &= \int_{(e_d=\bar{e})} f^*(\pi_d, x_d, e_d|\theta)\, d\pi_d dx_d \\ &= f_e^*(\bar{e})\,. \end{aligned}$$

2. Trading decisions $P_e(\tilde{e}|\pi_d, x_d, e_d)$ and $P_n(n|\pi_d, x_d, e_d)$ are defined in (1.11) and (1.7);

3. Offers and replies in the inter-dealer market achieves the optimum in (1.8) and (1.9);

4. Conjectures in (1.11), (1.7), (1.8), and (1.9) are correct given f^*;

5. Conjectures used for updating, in (1.5) and (1.4), are correct given f^*;

6. Letting $h^*(\theta)$ denote the long-run distribution of θ, and

$$f^*(\pi_d, x_d, e_d) = \sum_{\theta} f^*(\pi_d, x_d, e_d|\theta)\, h^*(\theta)\,,$$

$f^*(\pi_d, x_d, e_d|\theta)$ is the distribution of dealer's states (π_d, x_d, e_d) within the population implied by transitions (1.2), (1.5), (1.4), (1.3), and choice probabilities (1.11) and (1.7).[22]

1.5 Empirical Strategy

In this section, we lay out the empirical strategy followed to estimate the model described in Section 1.4. The main model primitives we wish to recover are dealers' trade costs, $\{c_0, c_1, c_{\text{buy}}, c_{\text{sell}}, c_{\tilde{e}}\}$, his inventory costs $\kappa(\cdot)$, as well as the standard deviations of cost shocks. We normalize β to 0.99.

The estimation proceeds in two steps. First, without directly leveraging the model, we recover dealers' experience $e_{d,t}$ and their beliefs $\pi_{d,t}$ about common preference shock θ_t. Next, we estimate dealers' choice probabilities and use them to recover the model primitives through Indirect Inference.

We begin by describing how we define and recover dealers' experience, as well as their beliefs. Next, we show how we recover dealers' search and inventory costs. Finally, in the Appendix, we describe how we recover the process for the unobserved preference shock θ. Results are presented in Section 1.6.

1.5.1 Experience

A sizable strand of literature has documented that more experienced firms have a competitive advantage in a variety of industrial settings (most notably, Benkard 2000, 2004). Experience is usually defined as the discounted sum of past production output. In turn, the marginal cost of production decreases as firms accumulate experience.

In this paper, we rely on the concept of experience to model dealers' incentives in the context of inter-dealer trade. We want to concisely capture the idea that dealers

[22]The steady state conditions are spelled out precisely Appendix 1.G.

select a trading counterparty based on the information about common shock θ_t that they will be able to extract from him. A dealer's experience offers a tractable way to proxy for the precision of the information that he has been able to gather through trade. In this Section we describe how we define experience in the data and, therefore, how we parametrize and recover experience's transition matrix r.

Dealers gain experience both by trading with retail investors and by trading with one another. The information content of inter-dealer trade will depend on how informed the trading counterparty is. For this reason, in contrast to the literature on learning-by-doing, it is not sufficient to keep track the sheer number of trades, but we also need to account for the experience of the trading counterparty.

Concretely, let g_t denote the network defined by inter-dealer trade during month t, with generic entry $g_{d,\tilde{d},t} = \mathbb{I}\left\{d \text{ and } \tilde{d} \text{ traded in } t\right\}$ and generic row $g_{d,t}$. Moreover, let $|n_{d,t}|$ denote the total quantity traded by dealer d with private investors in period t. We assume that the experience of dealer d, at the end of month t, for a given asset[23] is defined as

$$e_{d,t} = \delta e_{d,t-1} + |n_{d,t}| + \delta\psi_0 g_{d,t}\left(e_{t-1} - \delta e_{t-2}\right), \tag{1.12}$$

with initial values $e_{d,0} = 0$ and $e_{d,-1} = 0$.[24]

In Equation 1.12, parameter δ captures the idea that information becomes less relevant over *time*. Parameter ψ_0 captures the idea that the quality of a piece of information decays every time it is repeated to another agent.

[23]For tractability we cluster the assets traded in our sample into 15 groups, and estimate the experience process independently across groups.

[24]We experiment with numerous formulations for experience and find similar results. Among others, we tries: weighting $g_{d,t}$ by the size of trades; using the total volume of trade, rather than $|n_{d,t}|$; using the logarithm of total volume of trade, rather than $|n_{d,t}|$. Finally, definition 1.12 implicitly assumes that dealers only communicate the information they acquired in the previous period. This assumption can be easily relaxed without drastically changing the results. The advantage of the current set-up is that it minimizes double counting.

Lemma 1.5.1 shows that dealer d's experience in period t, $e_{d,t}$, can be rewritten as

$$e_{d,t} = \sum_{k \geq 0} \delta^k r_{d,t-k}, \tag{1.13}$$

where

$$\begin{aligned} r_{d,t} &= |n_{d,t}| + \delta\psi_0 g_{d,t} |n_{t-1}| + \delta\psi_0 g_{d,t} g_{t-1} |n_{t-2}| + \dots \\ &= |n_{d,t}| + \sum_{k \geq 1} (\delta\psi_0)^k g_{d,t} g_{t-1} \dots |n_{t-k}| . \end{aligned} \tag{1.14}$$

This rewriting is useful for interpreting Equation 1.12. Intuitively, $r_{d,t}$ describes the amount of information obtained by dealer d in period t. The first term in 1.14, $|n_{d,t}|$, captures the information obtained by dealer d directly by trading with investors. In period t dealer d also obtains information through inter-dealer trade. First, his direct counterparties will share some of the information that they acquired in the previous period by trading with investors. This is captured by the second term in 1.14, $\delta\psi_0 g_{d,t} |n_{t-1}|$. Dealer d's counterparties will also share some of the information that they acquired from *their* trading partners (and so on...), as captured by terms $(\delta\psi_0)^2 g_{d,t} g_{t-1} |n_{t-2}|$, $(\delta\psi_0)^3 g_{d,t} g_{t-1} g_{t-2} |n_{t-3}|$,.... In sum, dealer d's experience captures in a stylized way the discounted amount of information that that dealer d has obtained up to period t as a result of trade with dealers and investors. Let $r_t = e_t - \delta e_{t-1}$, then r_t satisfies

$$r_t = n_t + \sum_{k \geq 1} (\delta\psi_0)^k g_t g_{t-1} \dots n_{t-k}.$$

Next, we show that experience is bounded. Let D be the number of dealers in the market, and assume that $|n_t| \leq N$. If $\delta^2\psi_0 < \frac{1}{DN}$, for every g_t and n_t, the experience process $e_t = (e_{1t}, \cdots, e_{D,t})$ is bounded.

Estimation Strategy. Dealer's experience, as defined in (1.12), depends on two parameters: δ captures the rate of depreciation of information over time, while ψ_0 captures frictions that hinder communication in the context of inter-dealer trade. Due to data limitations, we normalize ψ_0 and focus on the estimation of the depreciation rate δ.

To estimate the parameters, we use a key implication of the model described in Section 1.4: trading with more experienced counterparties allows the dealers to observe a more informative signal about the state of the market. For this reason, sellers will be willing to charge a lower price to trade with a more experienced counterparty.

We draw from the literature on learning-by-doing, Benkard 2000 especially, to bring this implication of the model to the data. We consider prices p_i in transactions in the inter-dealer market, and estimate baseline specification

$$\log(p_i) = \alpha_{s_i \times \mathrm{m}_i \times a_i} + \alpha_{b_i} + \psi_1 \log\left(e_{b_i,\mathrm{m}_i,a_i}(\delta)\right) + \psi_2 x_{s_i,\mathrm{m}_i} + u_i, \tag{1.15}$$

where s_i, b_i, a_i and m_i denote, respectively, the seller, buyer, the asset involved in the trade, and the month in which trade happens. The parameter of interest in Equation 1.15 is δ. We include the coefficient ψ_1 to translate the units of the experience measure into dollars. This parameter measures the discount that a seller is willing to grant to an experienced buyer, and can be thought as a reduced-form measure of a dealers' value for information.[25]

Identification of the parameters Equation 1.15 relies on comparisons of inter-dealer prices in trades for specific asset a_i, seller s_i and month m_i. Especially, Equation 1.15 attributes systematic differences in price across trades executed by seller s_i in

[25]This parameter doesn't have a structural interpretation and we don't use it anywhere else in the estimation. Nevertheless, the sign and magnitude of the estimates of this parameter, reported in Section 1.6, provide further reduced form evidence about the relevance of experimentation motives for trade.

month m_i to differences in experience level $e_{b_i,\mathrm{m}_i,a_i}(\delta)$ of the buyers involved in the transactions. The fixed effect $\alpha_{s_i\times\mathrm{m}_i\times a_i}$ absorbs market-wide shocks to prices, as well as the seller's persistent heterogeneity that might affect prices. We also control for the seller's inventory in x_{s_i,m_i}.

To estimate Equation 1.15 we exploit the Generalized Method of Moments. The persistency of buyers' experience, together with the inclusion of fixed effect α_{b_i}, raises concerns in the spirit of Arellano and Bond 1991. For this reason we use the lagged values of volume of trade, $n_{b_i,\mathrm{m}_i-k,a_i}$, and centrality, $c_{b_i,\mathrm{m}_i-k,a_i}$, as instruments.

Endogeneity of u_i could be a potential concern. One could especially worry about persistent (but non-constant) and unobserved heterogeneity that allows buyers both to strengthen their trading activity and to pay lower prices in the inter-dealer market. To control for this scenario, similar to Li and Schürhoff 2014, we also estimate Equation 1.15 using the monthly aggregate municipal bond mutual fund outflows and inflows, as well as entry and exit of dealers from the market as instruments for buyers' experience, $e_{b_i,\mathrm{m}_i,a_i}(\delta)$. Identification of the parameters, in this case, relies on the fact that different dealers have different exposure to shocks captured by the instruments, due to their inventory or clientele.[26]

1.5.2 Dealers' Information

Dealers' beliefs $\pi_{d,t} \in \Delta(\Theta)$ about the current value of the unobserved common preference shock θ_t are a key variable to predict their choices. Traditionally, estimating models where learning is explicitly accounted for requires cumbersome computational methods to simulate and integrate out the players' unobservable beliefs. However, thanks to the assumptions made in Section 1.4 as well as to the granularity of our data set, we can recover dealers' beliefs $\vec{\pi}_t = (\pi_{d,t})_{d=1}^{D}$ for every period t.

[26] In light of the inclusion of the fixed effect $\alpha_{s_i\times\mathrm{m}_i\times a_i}$, we can ignore the impact of the shocks used as instrument on the market level of prices.

In particular, Proposition 1.5.2 shows that the updating rules 1.4 and 1.5 can be substantially simplified, as long as after trade the dealers communicate to one another their forecast for the prices. This allows us to recover dealers' beliefs without knowing the equilibrium strategies. Suppose that $y_{d,t} = \pi_{d,t}$, and let $f^*(y|\tilde{e}, \theta)$ denote the distribution of signal $\tilde{y}$ among agents with experience level $\tilde{e}$, conditional on common shock θ_t. The updating rules 1.4 and 1.5 become

$$\mathcal{L}_{\text{sell}}(\pi; \tilde{y}, \tilde{e})\left(\theta^k\right) = \frac{f^*\left(\tilde{y}|\tilde{e}, \theta^k\right)\pi\left(\theta^k\right)}{\sum_\theta f^*(\tilde{y}|\tilde{e}, \theta)\,\pi(\theta)}, \tag{1.16}$$

and

$$\mathcal{L}_{\text{buy}}(\pi; \tilde{y}, \tilde{e})\left(\theta^k\right) = \frac{f^*\left(\tilde{y}|\tilde{e}, \theta^k\right)\pi\left(\theta^k\right)}{\sum_\theta f^*(\tilde{y}|\tilde{e}, \theta)\,\pi(\theta)}. \tag{1.17}$$

We use updating rule 1.3, together with 1.16 and 1.17 to recover dealers' beliefs. This requires running a fixed-point algorithm to recover distribution f^*. Indeed, these updating rules depend on the distribution of dealers' beliefs f^*. In turn, the distribution used for updating affects the evolution of dealers' beliefs.

In brief, our algorithm proceeds in the following steps:

1. Initialize $\left(\vec{\pi}_t^{(0)}\right)_{t\geq 1}$ with $\vec{\pi}_t^{(0)} \equiv \bar{\pi}$, where $\bar{\pi}$ satisfies $\bar{\pi} = h\bar{\pi}$
2. Given a guess for dealers' beliefs $\left(\pi_t^{(m)}\right)_{t\geq 1}$, compute the distribution of dealers' beliefs after inter-dealer trade
   (a) First, given beliefs $\left(\pi_t^{(m)}\right)_{t\geq 1}$, compute the distribution of dealers' beliefs after trade with investors, conditional on experience and the unobserved state θ. In particular, for every d and t, compute $\hat{\pi}_{d,t}^{(m)} = \mathcal{L}_{inv}\left(h'\pi_{d,t}^{(m)}; \vec{p}_{n,t}\right)$, and estimate $\hat{f}^{(m)}\left(\hat{\pi}_{d,t}^{(m)}|\theta^k, e_{d,t}\right)$.

(b) Next, update dealers' beliefs based on the interaction on the inter-dealer market. Dealer d, with beliefs $\hat{\pi}_{d,t}^{(m)}$, after buying an asset from dealer $\tilde{d}$ updates according to

$$\hat{\hat{\pi}}_{d,t}^{(m)} = \mathcal{L}_{\text{buy}}\left(\hat{\pi}_{d,t}^{(m)}; y_{\tilde{d},t}, e_{\tilde{d},t}\right).$$

For sellers instead

$$\hat{\hat{\pi}}_{d,t}^{(m)} = \mathcal{L}_{\text{sell}}\left(\hat{\pi}_{d,t}^{(m)}; y_{\tilde{d},t}, e_{\tilde{d},t}\right).$$

(c) set $\pi_{d,t}^{(m+1)} = h'\hat{\hat{\pi}}_{d,t}^{(m)}$.

3. If $\int \left|\hat{f}_0^{(m+1)}\left(\hat{\pi}_{d,t}^{(m)}|\theta_t, e_{d,t}\right) - \hat{f}_0^{(m-1)}\left(\hat{\pi}_{d,t}^{(m)}|\theta_t, e_{d,t}\right)\right| \hat{f}_0^{(m-1)}\left(\hat{\pi}_{d,t}^{(m)}|\theta_t, e_{d,t}\right) d\pi < \epsilon$, set

$$(\pi_t^*)_{t\geq 1} = \left(\pi_t^{(m)}\right)_{t\geq 1}.$$

Otherwise repeat steps 2 and 3.[27]

1.5.3 Dealers' Costs

We next turn to the model's primitives: inventory cost $\kappa = \{\kappa_0, \kappa_1\}$, costs to trade with investors $\{c_{\text{buy}}, c_{\text{sell}}, c_1, c_2\}$, costs to trade with other dealers $\{c_{\tilde{e}}\}_{\tilde{e}=1}^{E}$ and the standard deviations of cost shocks $\{\sigma_\xi, \sigma_0, \sigma_1\}$. To obtain the parameters of interest, we use dealers' optimal trade choice probabilities defined in Equation 1.7 and 1.11, as well as prices on the inter-dealer market. Both of these are a function of the dealers' value functions, which in turn depend on the parameters.

[27] We assume that our data comes from one steady state, so that the distribution $f^*(\pi_{d,t}|\theta, e_{d,t})$ is fixed over time. Furthermore, the assumption of anonymity requires that dealers with the same experience are seen as equivalent from their peers, which implies that the update in (1.16) and (1.17) does not depend on d. Finally independence ensures that rules (1.16) and (1.17) do not depend on the dealer's own private state.

We estimate the parameters via Indirect Inference, following Gourieroux et al. 1993. As auxiliary models we employ an ordered Probit and multinomial Logit for, respectively, the trading decision with investors and trading decision with other dealers.

Concretely, we use the observed data to obtain the parameter estimates $\hat{\beta} = \left(\beta_{\text{aux}}, \left(\beta_{\text{aux}}^{\tilde{e}}\right)_{\tilde{e}=1}^{E}\right)$ and $\hat{\alpha} = \left(\alpha_{\text{aux}}^{n}\right)_{n=1}^{N}$ that maximize the likelihood associated to the auxiliary model

$$\begin{aligned}\mathcal{L}\left(\mathbb{I};\alpha,\beta,z\right) = &\sum_{d}\sum_{t}\sum_{n}\mathbb{I}_{d,t,n}\log\left(\Phi\left(z'_{d,t,n}\beta_{\text{aux}} - \alpha_{\text{aux}}^{n}\right) - \Phi\left(z'_{d,t,n}\beta_{\text{aux}} - \alpha_{\text{aux}}^{n-1}\right)\right)\\ &+\sum_{d}\sum_{t}\sum_{n}\mathbb{I}_{d,t,\tilde{e}}\log\left(\frac{\exp\left(z'_{d,t,n}\beta_{\text{aux}}^{\tilde{e}}\right)}{\sum_{\tilde{e}}\exp\left(z'_{d,t,n}\beta_{\text{aux}}^{\tilde{e}}\right)}\right),\end{aligned}$$

where $\mathbb{I}_{d,t,n}$ is an indicator equal to 1 if dealer d chooses n in period t; $\mathbb{I}_{d,t,\tilde{e}}$ is an indicator equal to 1 if dealer d sells to a dealer with experience $\tilde{e}$ in period t, and $z_{d,t,n}$ is the state variable of dealer d in period t, $z_{d,t} = (\pi_{d,t}, x_{d,t}, e_{d,t})$.

Next, at every guess of primitive parameter value $\tau = \left\{c_{\text{buy}}, c_{\text{sell}}, c_1, c_2, c_{\tilde{e}}, \kappa_0, \kappa_1, \sigma_0, \sigma_1, \sigma_\xi\right\}$, we use a nested fixed point algorithm to solve for the dealer's value functions (V, W) and generate M simulated data sets $\left\{\mathbb{I}_{d,t,\tilde{e}}^{(m)}, \mathbb{I}_{d,t,\tilde{e}}^{(m)}\right\}_{m=1}^{M}$. Finally, we use the simulated dataset to retrieve the parameters $\left(\beta^{(m)}(\tau), \alpha^{(m)}(\tau)\right)$ that maximize the auxiliary likelihood $\mathcal{L}\left(\tilde{\mathbb{I}};\alpha,\beta,z\right)$. The estimated primitive parameter $\hat{\tau}$ minimizes

$$\mathcal{L}\left(\mathbb{I};\hat{\alpha},\hat{\beta},z\right) - \mathcal{L}\left(\mathbb{I};\bar{\alpha}(\tau),\bar{\beta}(\tau),z\right),$$

where $\bar{\beta}(\tau) = \frac{1}{M}\sum_{m=1}^{M}\beta^{(m)}(\tau)$, and $\bar{\alpha} = \frac{1}{M}\sum_{m=1}^{M}\alpha^{(m)}(\tau)$. In Appendix 1.F we outline the specific steps we follow in the estimation algorithm.

Identification. Estimation of the parameters in τ relies on the dynamic panel nature of the data. Observing dealers' decision over time allows us to keep track

of how their behavior change depending on their type (π, x, e). These changes in behavior identify the parameters.

Conceptually, identification of standard deviations $\{\sigma_0, \sigma_1\}$ —associated to the cost shocks in the context of trade with investors— is driven by differences in dealers trading decisions across different beliefs π, conditional on observed trading prices. As standard deviation σ_0 —which affects the decision on whether to buy or sell the asset— increases, dealers will choose to either sell or purchase the asset with the same probability for any belief π. As standard deviation σ_1 increases, the dealer will randomize between trading the highest and the lowest possible amount of units of the asset. Finally, sheer trading costs $\{c_{\text{buy}}, c_{\text{sell}}, c_1, c_2\}$ are pinned down by the overall level of trade.

Identification of trade costs in inter-dealer trade $\{c_{\tilde{e}}\}_{\tilde{e}=1}^{E}$ depends in a similar fashion on dealers' decisions in the context of inter-dealer trade. Importantly, we rely on data on prices for inter-dealer trades to anchor the unit of utility from inter-dealer trade. This allows us to identify standard deviation σ_ξ .

Finally, prices in inter-dealer trades as a function of the seller's inventory x help us to identify inventory costs κ. If the inventory cost is high the trading price will sharply fall as dealer's inventory increases, holding everything else fixed.

1.6 Results

In this section we present the results from our empirical analysis. Throughout the estimation, we group assets into 15 different classes and consider each class as a separate market. The classes are determined, through a clustering algorithm, to maximize the correlation over time of the average price within each class.[28] Appendix 1.C describes how we recover Θ, sequence $(\theta_t)_{t\geq 1}$, and transition matrix h. For each class, common shock θ_t can take at most three values.

[28]For details, see Appendix 1.D

1.6.1 First stage

We now show the results from the first stage of the estimation described in Sections 1.5.1 and 1.5.2.

We estimate Equation 1.12 separately for each class of assets. Table 1.3 shows the average estimated parameters, across classes of assets, for the experience process. The top panel reports the results from our baseline estimation, while the bottom panel reports those from the instrumented version. Table 1.9 in the Appendix presents the results separately for each class.

For each class, the experience discount, as measured by ψ_1, is negative and significant at the 1% level. Dealers, therefore, are willing to pay a premium to trade with more experienced counterparties. The average trading price falls by 0.17 percentage points when the buyer belongs to the 75th percentile of the experience distribution, compared to when it belongs to the 25th percentile. This amounts to approximately 10% of the average intermediation spread.[29] Consistent with the volatile nature of the market, information appears to be short-lived. The value of δ is around 0.55 in both specifications. This implies that only 2% of the accumulated experience lives through six months.

To validate these results, in Figure 1.4 we compare the estimates of ψ_1 for each class of assets with the variability of prices over time for the same class of assets. Intuitively, for assets for which price is less uncertain, experience discount $|\psi_1|$ should be smaller. On the vertical axis of Figure 1.4 we plot the estimate for ψ_1. On the horizontal axis we report the R^2 from the regression of the current average price on market indexes. Assets for which the R^2 is larger are easier to forecast based on information publicly known. We find that for classes of assets associated with larger

[29] The learning rate $1 - 2^{\psi_1}$ captures the percentage reduction in price associated with a doubling of experience.

values of R^2, the experience discount is smaller. The (weighted) correlation[30] between the two values is 26%.

Table 1.3: Experience process

	Persistence of Information δ	Experience Premium ψ_1	Learning Rate $1-2^{\psi_1}$
I	0.582	-0.0017	0.118%
II	0.53	-0.0013	0.09%

Notes: The table above summarizes the estimates of the Experience process defined in Equation 1.12. We cluster the assets in our sample into 15 groups and estimate the experience process independently across groups. In the table above we report the average parameter across classes. For the top row we use past trading activity as an instrument. For the bottom row we use instruments defined in Section 1.5.1.

[30]We weight the correlation using the total value of the trades for each group of assets.

Figure 1.4: Comparison of risk and ψ_1 estimates across classes of assets.

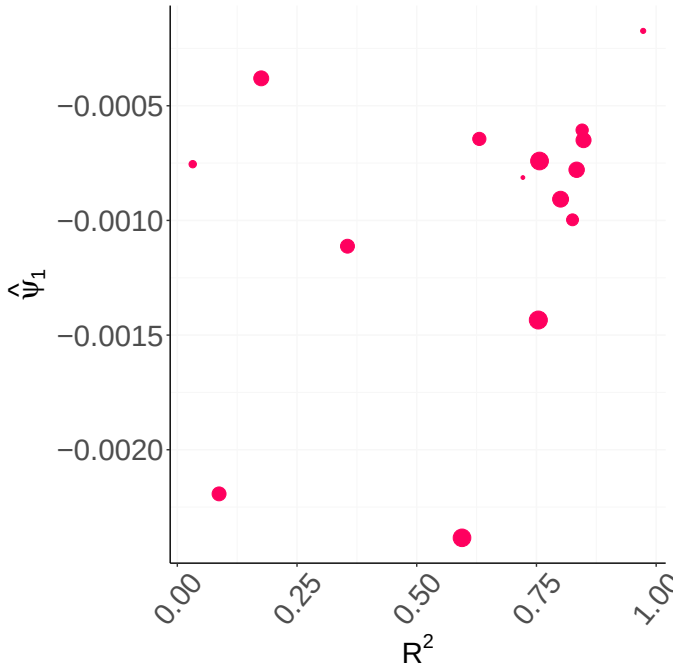

Notes: the above figure compares experience discount ψ_1 and price uncertainty across assets groups. On the vertical asset we plot the estimate of experience discount ψ_1. To proxy for the price uncertainty of different groups of assets, on the horizontal axis we plot the R^2 of a regression of monthly average price within a class of assets on public indexes about on the performance of the market for municipal bonds.

1.6.2 Dealers' Costs

Table 1.4 reports the average baseline estimates for the model primitives across classes of inventory. Namely, inventory costs $\{\kappa_0, \kappa_1\}$, costs to trade with investors $\{c_{\text{buy}}, c_{\text{sell}}, c_1, c_2\}$ and with other dealers $\{c_{\tilde{e}}\}_{\tilde{e}=1}^{E}$, and the standard deviations of cost shocks $\{\sigma_\xi, \sigma_0, \sigma_1\}$.

Trade costs are large, with the average dealer spending on average $\$3,000$ per class of assets each month to find investors. The total search cost is on average $\$50,000$ per month and dealer. The standard deviation of the preference shocks, σ_1, equals roughly 16% of the trading price, suggesting that that preference shocks ϵ^1 do not account for a disproportionately large part of the decision on how many investors to

trade with. Consistent with the industry narrative, dealers receive higher fees when they sell assets to investors, obtaining net fees of around $\$1,300$ per class of asset traded. Search costs, instead, dominate when it comes to buying assets, as dealers pay net fixed cost of around $\$1,000$ if they decide to buy the asset.

If we interpret the inventory cost in terms of leverage, the estimates imply that 15% of the inventory is collateralized, as long as the dealers can borrow at the deposit funds rates. The difference in cost of inventory across assets is explained by the rating of the assets traded. In particular, doubling the number of assets with a B-rating doubles the cost of inventory for a class of assets.[31] This is consistent with, as an example, dealers targeting a certain value-at-risk level when managing inventory.

Table 1.4: Baseline cost estimates

Trade with Investors		**Fees**		**Interdealer Trade**			**Inventory Costs**		**Variance**		
c_1	c_2	c_{buy}	c_{sell}	$c_{\tilde{e}_1}$	$c_{\tilde{e}_2}$	$c_{\tilde{e}_3}$	κ_0	κ_1	σ_ξ	σ_1	σ_0
0.018	0.012	0.023	-0.020	-0.040	-0.019	-0.020	0.003	0	0.11	0.19	0.10

Notes: The table above summarizes the estimates of the trading costs that dealer faces. We cluster the assets in our sample into 15 groups and estimate these costs independently across groups. In the table above we report the average parameter across classes.

1.7 Value and Precision of Information

Next, we explicitly characterize the dealers' incentives to experiment. We begin by comparing the informativeness of trades with investors and inter-dealer trades. Next, we show to what extent experimentation helps dealers refine their estimate about the state of fundamentals. Moreover, we study the origin of dealer's information. Finally,

[31] Rating explains 40% of the variation in inventory cost across classes.

we characterize the dealers' value of information. For all these exercises we use the biggest group of assets.

1.7.1 Precision of Information

To compare the informativeness of different types of signals, we study the precision of the forecast for the trading price $\mathbb{E}(p_i|\pi,\mathcal{I}) = \sum_{\theta\in\Theta}\pi(\theta|\mathcal{I})\,\mathbb{E}(p_i|\theta)$ of a dealer with belief π, who receives information $\mathcal{I}$. In particular we average (and square) the value of $[\mathbb{E}(p_i|\theta_t) - \mathbb{E}(p_i|\pi,\mathcal{I})]^2$ across different realizations of θ_t,

$$RMSE(\pi,\mathcal{I}) = \sqrt{\mathbb{E}\left([\mathbb{E}(p_i|\theta_t) - \mathbb{E}(p_i|\pi,\mathcal{I})]^2\right)}. \tag{1.18}$$

Intuitively, $\mathbb{E}\left([\mathbb{E}(p_i|\theta_t) - \mathbb{E}(p_i|\pi,\mathcal{I})]^2\right)$ captures the average difference in the mean squared error of the best prediction for trading price p_i of a dealer with belief $\pi(\theta|\mathcal{I})$, compared to a dealer who knew the realization of θ_t.

Figure 1.5 plots $RMSE(\pi,\mathcal{I})$ for different information sets $\mathcal{I}$ and dealers' prior beliefs π. In all cases $|\Theta| = 3$, so that $\pi \in \Delta^2$. We set the probability for the middle state to zero, and plot the probability of the low state, π_1, on the horizontal axes. Different lines correspond to different signals (i.e., different information sets $\mathcal{I}$). In the lower panel, we plot the informativeness of trade with investors for different numbers of trades. The upper panel depicts the informativeness of the signal obtained by trading with another dealer for different levels of experience. We also include the root mean squared error of the estimate of θ_t absent any additional information (orange line). As one could have guessed, the signal is more informative if the dealer trades more or if he trades with a more experienced dealer. Furthermore a dealer is more easily convinced if his prior beliefs are more dispersed. In particular, the root mean squared error falls at most by 40% if the dealer decides to trade with an experienced dealer (blue line), and at most by 20% if the dealer trades with an inexperienced

Figure 1.5: Informativeness of different signals

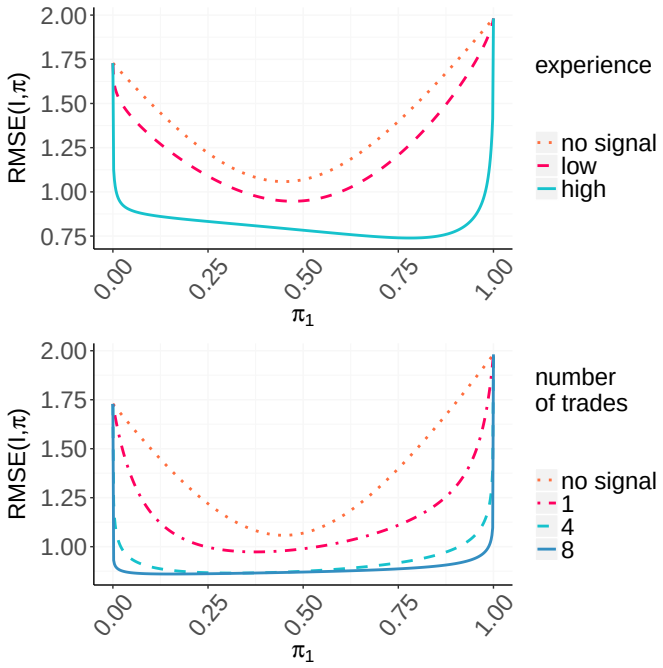

Notes: the above figure plots the root mean square error $RMSE(\pi, \mathcal{I})$ for different signals and dealers' prior beliefs π. In all cases $|\Theta| = 3$, and we set the probability for the middle state to zero. The probability of the low state π_1, is shown on the horizontal axes. Different lines correspond to different information sets $\mathcal{I}$. In the lower panel, we plot $RMSE(\pi, \mathcal{I})$ after dealers traded with investors. The upper panel depicts the value of $RMSE(\pi, \mathcal{I})$ after trading with another dealer.

counterparty. Trading with a single investor is as informative as trading with the average inexperienced trader, as the RMSE falls at most by 15% in the latter case. Trading with an experienced dealer instead is as informative as trading with 5 different investors.

Next, we look at the uncertainty in dealers' estimate of trading prices taking into account their equilibrium behavior. In Table 1.5 we compute the average RMSE across players, at the observed equilibrium, for different types of dealers. The first column reports the upper bound for the RMSE when the agents only observe public signal y_t^P. The last three columns show the RMSE attained in equilibrium for dealers with different experience levels. Experimentation allows dealers to improve the precision of

their estimate on average by 25%. Experienced dealers have better information than the rest of the market, as they are able to improve the precision of their prediction by 33% compared to when they did not experiment. The improvement in precision is lower for inexperienced dealers, who improve the precision of their estimate by 23%.

Table 1.5: Precision of information

	Uninformed Players	**Market Average**	**Inexperienced Players**	**Experienced Players**
RMSE	1.024	0.759	0.782	0.691
Percentage	100	74.15	76.43	67.49

Notes: We compute the measure defined in Equation 1.18 for dealers along the equailibrium path. The table above reports the average across beliefs π for different classes of dealers. The first column reports the measure for players with access only to public information. The second column reports the average across all players, and the last two columns distinguish among experienced and inexperienced players.

Next we study how information percolates in the market. To this end, we use the estimated policy functions to simulate the game under the assumption that dealers only update based on information from inter-dealer trade with a specific type of dealer or based on trade with investors. Table 1.6 reports the percentage of the increase in precision that can be attributed to different sources of information. Experienced dealers learn mainly from other dealers, as only 19% of their information derives from trade with investors. Inexperienced dealers rely more heavily on trade with investors, which accounts for one-third of their information. For both experienced and inexperienced players, trade with inexperienced dealers accounts for the largest share of information acquired the context on inter-dealer trade. The difference is sharper for inexperienced dealers: only 22% of information that inexperienced dealers gather derives from trade with experienced dealers.

Table 1.6: Origin of information

	Trade with Investors	**Inter-dealer Trade**		
		Low Experience	Middle Experience	High Experience
Experienced Players	0.19	0.288	0.262	0.254
Inexperienced Players	0.304	0.30	0.278	0.223

Notes: The table above decomposes improvement in dealers' information described in Table 1.5. The first column shows what percentage of dealers' information derives from trade with investors. The last three columns show what percentage of dealers' information derives from trade with dealers with different experience levels.

1.7.2 Value of Information

To quantify dealers' incentives to experiment, we characterize the value of information for the dealers active in the market for municipal bonds. Intuitively, in our estimation the dealers' value for information is identified using prices in the inter-dealer market: the discount that a dealer is willing to offer to trade with a more experienced counterparty allows us to measure the value he assigns to a more informative signal.

For a dealer with prior beliefs π, inventory x and experience level e, the value of a signal $y|\theta \sim f_y$ is

$$VI\left(\pi, x, e, f_y\right) = \mathbb{E}\left(V\left(\mathcal{L}\left(y, \pi\right), x, e\right)\right) - V\left(\pi, x, e\right), \tag{1.19}$$

where the first expectation is taken with respect to realizations of y, and $\mathcal{L}\left(y, \pi\right)$ denotes the updated belief based on observing the realization of signal y:

$$\begin{aligned}\pi'\left(\theta^k\right)\left(y\right) &= \frac{f\left(y|\theta^k\right)\pi\left(\theta^k\right)}{\sum_{\theta\in\Theta} f\left(y|\theta\right)\pi\left(\theta\right)} \\ &= \mathcal{L}\left(y, \pi\right).\end{aligned}$$

The value of VI can be interpreted as the highest price that a dealer with beliefs π, inventory x and experience level e is willing to pay to purchase signal y about θ. The upper panel of Figure 1.6 shows the value of information for different prior beliefs π and different distributions f_y. For simplicity, we focus on normal signals, $y|\theta \sim N(\theta, \sigma_I^2)$, and normalize the precision of signal y, $\frac{1}{\sigma_I}$, based on the standard deviation of the public signal about θ, $\frac{1}{\sigma_I} = \frac{\gamma}{\sigma_P}$. Finally, to fix magnitudes, we normalize value of information using the average intermediation spread. In sum, the plot should be interpreted as the share of the intermediation spread that a dealer is willing to give up to buy a signal about θ which is γ times more informative than the public signal y_t^P.

The lower panel of Figure 1.6 shows the marginal value of information

$$MVI(\pi, x, e, f_y) = \frac{VI(\pi, x, e, f_{y_h}) - VI(\pi, x, e, f_y)}{h}, \tag{1.20}$$

where $f_{y_h} = N\left(\theta, \left(\frac{\sigma_P}{\gamma+h}\right)^2\right)$ is the distribution of a slightly more precise signal than $f_y = N\left(\theta, \left(\frac{\sigma_P}{\gamma}\right)^2\right)$.

Both value and marginal value of information are computed at the average level of inventory $\bar{x}$ and normalized using the average intermediation spread.

The estimated value of information is positive and substantial. Dealers benefit from having precise information about the market fundamental θ. As an example, the average dealer is willing pay up to 15% of its intermediation spread to acquire a signal about θ twice as informative as the public signal.

The marginal value of information is initially zero as found in Radner and Stiglitz 1984 and rigorously formalized in Chade and Schlee 2002. The marginal value is hump shaped. This is consistent with what Keppo et al. 2008 find in a Gaussian setting. The shape of the marginal value can have far-reaching implications for the

market. For example, it creates a scope for a dealer's specialization that is consistent with the tendency of dealers to specialize in a specific class of assets.

Figure 1.6: Value and marginal value of information

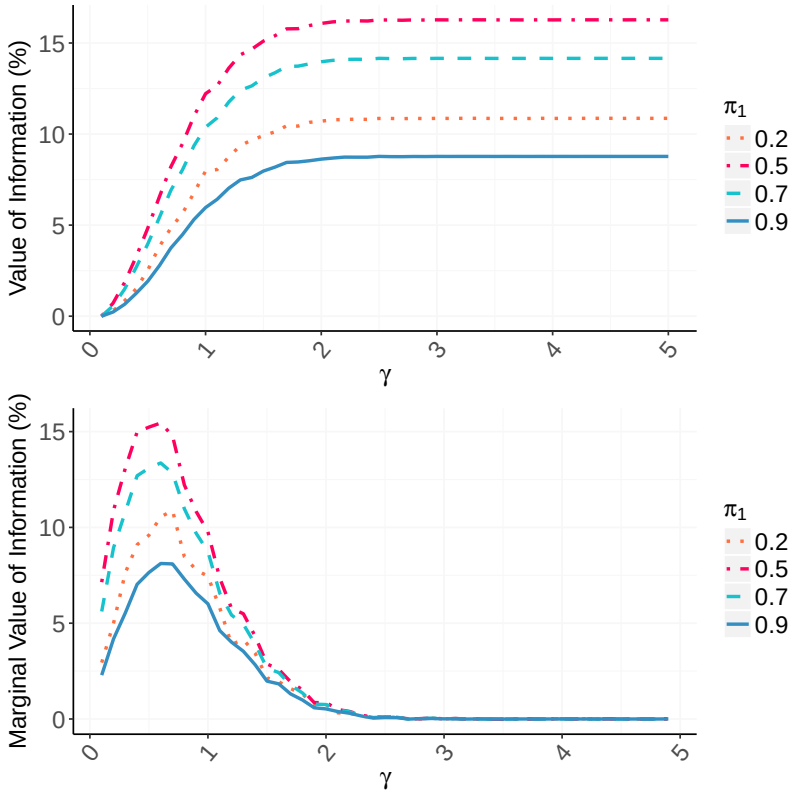

Notes: The top panel of the above figure plots the share of the average intermediation spread that a dealer is willing to give up to buy a noisy signal of common shock θ with distribution $y|\theta \sim N\left(\theta, \frac{\gamma}{\sigma_P}\right)$, where σ_P is the variance of the public signal y_t^P. Different lines corresponds to different prior beliefs π. In particular, in all cases $|\Theta| = 3$, and we set the probability for the middle state to zero. Each line correspond to a different probability of the low state π_1. Instead, the bottom panel of the figure plots the difference in the dealer's valuation for a signal $y|\theta \sim N\left(\theta, \frac{\gamma}{\sigma_P}\right)$ and a slightly less informative signal $N\left(\theta, \frac{\gamma-h}{\sigma_P}\right)$ normalized by h, for small values of h.

1.8 Implications for Market Transparency

Traditionally, assets in decentralized markets are traded in an opaque environment, with limited or no public information about market activity. In the last decade, however, access to trade information has improved in many decentralized markets, mostly due to direct intervention of the policy maker. As an example, in 1995 the MSRB took the first steps of the plan that led, in 2005, to the 15-minute reporting discussed in Section 1.3. Shortly after, in 2002, a similar provision was imposed by The Financial Industry Regulatory Authority (FINRA) in the market for U.S. corporate bonds. In 2011 Agency-Backed Securities and Asset-Backed Securities followed. Finally, in 2014 FINRA began disseminating 144A transactions.

The push toward greater transparency in decentralized markets is still ongoing, both in the US and abroad. As an example, in July FINRA began requiring its member firms to report U.S. Treasury securities transactions, even though those prices are currently not disseminated to the public. In Europe the legislative package comprising the revised Markets in Financial Instruments Directive and a new Regulation ("MiFID II"), passed into law in 2014, is about to institute a post-trade transparency regime which will affect a broad range of instruments.

The stated objective behind these policies is to increase the assets' liquidity by improving investor participation and trade activity.[32] An asset is considered more liquid "if it is more certainly realizable at short notice without loss." Therefore, liquidity is valuable per se, as long as investors value immediacy. Moreover, a liquid secondary market is a crucial condition to lower the cost of raising capital. As an example, Wang et al. 2008 estimate that the municipal bond issuers pay 13 billion a year to compensate investors for the risks implied by the illiquidity of the market.

[32]Testimony of Chairman Arthur Levitt Before the House Subcommittee on Finance and Hazardous Materials, Committee on Commerce, Concerning Transparency in the United States Debt Market and Mutual Fund Fees and Expenses.

Increasing liquidity in this market, therefore, would translate to huge savings for local governments and municipalities.

The argument for the effectiveness of the policy goes as follows: dealers have an informational advantage vis-a-vis investors. This advantage is leveraged to "buy low and sell high." This lowers liquidity directly, since it lowers the price that an investor can obtain by selling his assets before maturity. Moreover, dealers' market power lowers liquidity indirectly, since higher prices on the buy side depress buyers' participation. Improving access to public information, therefore, should reduce dealers' market power and increase market liquidity.

This argument, however, ignores a key driving force of market liquidity: the dealers' incentives to participate to trade. In particular, if information acquisition is a key determinant of dealers' trading decisions, improving access to public information might weaken their incentives to trade and dampen or overturn the positive effects of investors' participation.

We use the estimated model to qualify this statement. In particular, we quantify the effect of an increase in market transparency of dealers' incentives to trade. To approximate a transparent market, we simulate the model assuming that the terms of trade of all transactions become public at the end of each period. Once public, information about trade activity can be observed, free of charge, by everyone.

On average dealers are willing to buy 4% fewer assets from investors, once market transparency is improved. In particular, Table 1.7 shows that on average (across assets) dealers' purchase increase by 4% when the market value for the asset increase. This change is offset by the change in trading behavior for other realizations of the asset.

Two effects are at play. First, transparency weakens the incentives to experiment: when information trading activity is made public, uncertainty about common shock θ_t is drastically reduced. Therefore, the value of additional information conveyed by

Table 1.7: Effect of market transparency

State (θ_t)	**Low**	**Medium**	**High**
Purchases from investors (Overall Change)	4.97%	−4.72%	−8.68%
(Billion $)	8.00	−5.1	−14.95

Notes: the table above summarizes the change in volume of dealers' assets purchases that results from an increase in market transparency. In particular we report the average change across different classes of assets, for different values of common preference shock θ_t.

trade becomes irrelevant. This makes each trade less valuable for the dealers and implies that, conditional on private history (π, x, e), dealers are willing to trade more sporadically. Second, improving public information reduces uncertainty about the realization of the preference shock θ_t. Lower uncertainty implies that dealers are more willing to trade larger quantities of the asset, partially offsetting the first effect.

The balance between these two effects varies substantially across assets. As shown in Figure 1.7, the average change within classes of assets ranges from -10% to $+10\%$. This suggests that the success of this policy will hinge on the underlying features of the assets traded.

Estimated cost of inventory $\hat{\kappa}_0$ is the strongest predictor for the effect of market transparency, and accounts for more that 50% of the differences in market outcome across classes of assets. This result is intuitive: as the limits to expand inventory become stronger, the dealer will find it more difficult to adjust his trading decisions to exploit fluctuations in common shock θ_t. This contains the first effect described above, since the dealer is in a worse position to leverage information acquired thanks to market transparency, and it implies that the decline in volume of trade is sharper.

Figure 1.7: Total change in volume of trade as a function of cost of inventory

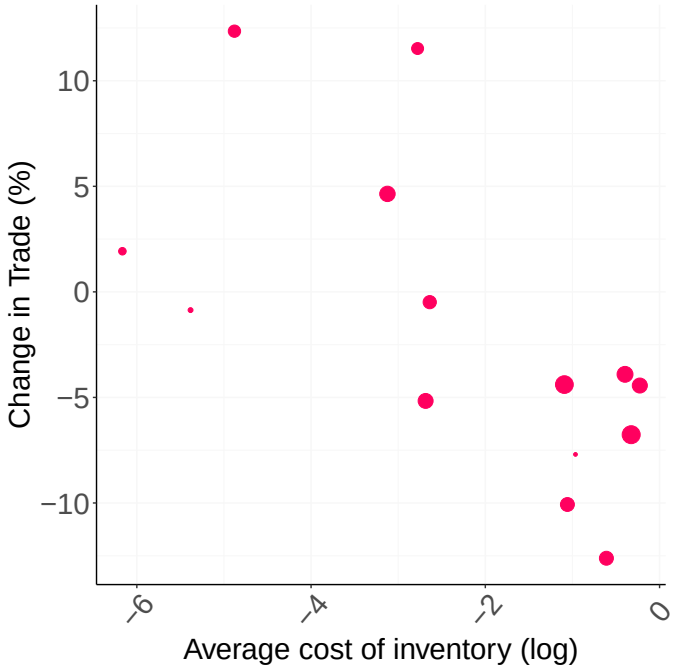

Notes: The figure above compares the total change in volume associated to an increase in market transparency with the estimated cost of inventory κ_0 across different classes of assets. For each class of assets, the size of the associated point is proportional to the observed total value of trade for the assets included in the class.

1.9 Conclusion

In this paper, we shed new light on the role of experimentation in decentralized opaque markets. These markets are common in wholesale trade markets and markets for investment goods. We argue that in these markets trade can be a source of valuable information about the market fundamentals. Obtaining this information, therefore, becomes an additional motive for trade.

To characterize incentives to experiment, we use a detailed dataset of transactions on the secondary market of municipal bonds, which provides a comprehensive insight into a decentralized financial market. We first use the dataset to provide reduced form evidence suggesting that incentives to experiment are a first-order motive for trade in the market. To rationalize these facts we build a dynamic model of trade in decentralized markets where agents are uncertain about the underlying value of the asset traded. Using the data we estimate the model and demonstrate that experimentation explains up to 10% of the volume of trade in this market. Finally we show that accounting for experimentation is important for a number of policies, such as increasing market transparency as well as imposing limits on inventory holding.

References

Admati, A. R. and Pfleiderer, P.: 1988, A theory of intraday patterns: Volume and price variability, *The Review of Financial Studies* **1**(1), 3–40.

Aghion, P., Espinosa, M. P. and Jullien, B.: 1993, Dynamic duopoly with learning through market experimentation, *Economic Theory* **3**(3), 517–539.

Aguirregabiria, V. and Mira, P.: 2007, Sequential estimation of dynamic discrete games, *Econometrica* **75**(1), 1–53.

Arellano, M. and Bond, S.: 1991, Some tests of specification for panel data: Monte carlo evidence and an application to employment equations, *The Review of Economic Studies* **58**(2), 277–297.

Asquith, P., Covert, T. and Pathak, P.: 2013, The effects of mandatory transparency in financial market design: Evidence from the corporate bond market, *National Bureau of Economic Research* .

Bajari, P., Benkard, C. L. and Levin, J.: 2007, Estimating dynamic models of imperfect competition, *Econometrica* **75**(5), 1331–1370.

Benkard, C. L.: 2000, Learning and forgetting: The dynamics of aircraft production., *American Economic Review* **90**(4), 1034–1054.

Benkard, C. L.: 2004, A dynamic analysis of the market for wide-bodied commercial aircraft, *The Review of Economic Studies* **7**(31), 581–611.

Bloomfield, R. and O'Hara, M.: 1999, Market transparency: Who wins and who loses?, *The Review of Financial Studies* **12**(1), 5–35.

Bloomfield, R. and O'Hara, M.: 2000, Can transparent markets survive?, *Journal of Financial Economics* **55**(3), 425–459.

Blouin, M. R. and Serrano, R.: 2001, A decentralized market with common values uncertainty: Non-steady states, *The Review of Economic Studies* **68**(2), 323–346.

Chade, H. and Schlee, E.: 2002, Another look at the radner–stiglitz nonconcavity in the value of information, *Journal of Economic Theory* **107**(2), 421–452.

Chandrasekhar, A. G., Larreguy, H. and Xandri, J. P.: 2012, Testing models of social learning on networks: Evidence from a framed field experiment, *Work. Pap., Mass. Inst. Technol., Cambridge, MA* .

Duffie, D., Gârleanu, N. and Pedersen, L. H.: 2005, Over-the-counter markets, *Econometrica* **73**(6), 1815–1847.

Duffie, D., Gârleanu, N. and Pedersen, L. H.: 2007, Valuation in over-the-counter markets, *The Review of Financial Studies* **20**(6), 1865–1900.

Duffie, D. and Manso, G.: 2007, Information percolation in large markets, *The American economic review* **97**(2), 203–209.

Ellison, G. and Fudenberg, D.: 1993, Rules of thumb for social learning, *Journal of Political Economy* **101**(4), 612–643.

Ellison, G. and Fudenberg, D.: 1995, Word-of-mouth communication and social learning, *The Quarterly Journal of Economics* **110**(1), 93–125.

Ericson, R. and Pakes, A.: 1995, Markov-perfect industry dynamics: A framework for empirical work, *The Review of Economic Studies* **62**(1), 53–82.

Farboodi, M., Jarosch, G. and Shimer, R.: 2016, Meeting technologies in decentralized asset markets, *Working Paper* .

Gavazza, A.: 2011a, Leasing and secondary markets: Theory and evidence from commercial aircraft, *Journal of Political Economy* **119**(2), 325–377.

Gavazza, A.: 2011b, The role of trading frictions in real asset markets, *American Economic Review* **101**(4), 1106–43.

Golosov, M., Lorenzoni, G. and Tsyvinski, A.: 2014, Decentralized trading with private information, *Econometrica* **82**(3), 1055–1091.

Gourieroux, C., Monfort, A. and Renault, E.: 1993, Indirect inference, *Journal of Applied Econometrics* **8**(S1).

Green, R. C., Li, D. and Schürhoff, N.: 2010, Price discovery in illiquid markets: Do financial asset prices rise faster than they fall?, *The Journal of Finance* **65**(5), 1669–1702.

Grossman, S. J., Kihlstrom, R. E. and Mirman, L. J.: 1977, A bayesian approach to the production of information and learning by doing, *The Review of Economic Studies* pp. 533–547.

Harris, L. E. and Piwowar, M. S.: 2006, Secondary trading costs in the municipal bond market, *The Journal of Finance* **61**(3), 1361–1397.

Hopenhayn, H. A.: 1992, Entry, exit, and firm dynamics in long run equilibrium, *Econometrica* **60**(5), 1127–1150.

Hörner, J. and Skrzypacz, A.: 2016, Learning, experimentation and information design.

Hugonnier, J., Lester, B. and Weill, P.-O.: 2014, Heterogeneity in decentralized asset markets, *National Bureau of Economic Research* .

Kalouptsidi, M.: 2014, Detection and impact of industrial subsidies: The case of world shipbuilding, *American Economic Review* **104**(2), 564–608.

Keppo, J., Moscarini, G. and Smith, L.: 2008, The demand for information: More heat than light, *Journal of Economic Theory* **138**(1), 21–50.

Kihlstrom, R., Mirman, L. and Postlewaite, A.: 1984, Experimental consumption and the 'rothschild effect', *Bayesian Models of Economic Theory* .

Kyle, A. S.: 1985, Continuous auctions and insider trading, *Econometrica* pp. 1315–1335.

Kyle, A. S.: 1989, Informed speculation with imperfect competition, *The Review of Economic Studies* **56**(3), 317–355.

Leach, J. C. and Madhavan, A. N.: 1993, Price experimentation and security market structure, *The Review of Financial Studies* **6**(2), 375–404.

Li, D. and Schürhoff, N.: 2014, Dealer networks.

Mirman, L. J., Samuelson, L. and Urbano, A.: 1993, Monopoly experimentation, *International Economic Review* pp. 549–563.

Pagano, M.: 1989, Trading volume and asset liquidity, *The Quarterly Journal of Economics* **104**(2), 255–274.

Pakes, A., Ostrovsky, M. and Berry, S. T.: 2007, Simple estimators for the parameters of discrete dynamics games, *The RAND Journal of Economics* **38**(2), 373–399.

Radner, R. and Stiglitz, J.: 1984, A nonconcavity in the value of information, *Bayesian models in Economic Theory* **5**, 33–52.

Rust, J.: 1987, Optimal replacement of gmc bus engines: An empirical model of harold zurcher, *Econometrica* **55**(5), 999–1033.

Schultz, P.: 2012, The market for new issues of municipal bonds: The roles of transparency and limited access to retail investors, *Journal of Financial Economics* **106**(3), 492–512.

Wang, J., Wu, C. and Zhang, F. X.: 2008, Liquidity, default, taxes, and yields on municipal bonds, *Journal of Banking & Finance* **32**(6), 1133–1149.

Weintraub, G. Y., Benkard, C. L. and Van Roy, B.: 2008, Markov perfect industry dynamics with many firms, *Econometrica* **76**(6), 1375–1411.

Wolinsky, A.: 1990, Information revelation in a market with pairwise meetings, *Econometrica* pp. 1–23.

Appendix

1.A Additional Tables and Figures

	Sales Trade Size			Purchases Trade Size		
	6 Months	9 Months	12 Months	6 Months	9 Months	12 Months
uninsured * $\mathbb{I}\{t > t_0\}$	−6.145	−2.346	−1.665	−5.213	−7.674	−6.713
	(7.222)	(3.898)	(3.297)	(4.420)	(5.946)	(5.054)
N	289,886	476,903	640,087	320,509	429,568	451,482
Level	Issuer-Week					

Table 1.8: Difference-in-Difference Estimates of Transparency on Trade Size

	Baseline Specification		Instrumented Version	
Group	β	δ	β	δ
I	-0.0008	0.849	-0.0002	0.249
II	-0.0116	0.948	-0.0003	0.634
III	-0.0005	0.763	0.00012	0.601
IV	-0.0006	0.829	-0.0001	0.667
V	-0.0001	0.903	-0.0002	0.236
VI	-0.0023	0.616	-0.0007	0.288
VII	-0.0003	0.255	-0.0002	0.252
VIII	-0.0006	0.257	-0.0004	0.752
IX	-0.0007	0.321	-7.92e-	0.815
X	-0.0016	0.572	-0.0006	0.255
XI	-0.0007	0.348	-0.0007	0.274
XII	-0.0008	0.232	-0.0010	0.397
XIII	-0.0005	0.964	-0.0137	0.887
XIV	-0.0006	0.276	-0.0008	0.969
XV	-0.0025	0.368	-0.0010	0.958
XVI	-0.0010	0.791	-0.0016	0.289

Table 1.9: Estimates for Experience

1.B What do Dealers Know?

We use the specification test suggested in Dickstein and Morales (2015) to test the assumption that dealers have no information about the market value of the asset in months where they don't participate to trade. The intuition behind the test is the following: let $y_{d,t}$ be an outcome variable that depends on a decision of dealer d in period t, such as the quantity traded in a certain asset, or the price charged to investors. Let also $\mathcal{I}_{d,t}$ denote dealer d's information set at the beginning of period t. Dealer d's decision about $y_{d,t}$ will depends on dealer d's expectation of the market value for the asset $\mathbb{E}\left(\theta_t|\mathcal{I}_{d,t}\right)$, conditional on what he knows about past realizations of θ_t. Under this scenario, if a variable Z_t belongs to d's information set $\mathcal{I}_{d,t}$, then it must be orthogonal to his forecast error:

$$\mathbb{E}\left[\left(\theta_t - \mathbb{E}\left(\theta_t|\mathcal{I}_{d,t}\right)\right) Z_t\right] = 0.$$

In this case, therefore, Z_t would be a instrument for $\mathbb{E}(\theta_t|\mathcal{I}_{d,t})$ in the regression

$$\begin{aligned} y_{d,t} &= \alpha + \beta\theta_t + \beta\left(\mathbb{E}(\theta_t|\mathcal{I}_{d,t}) - \theta_t\right) \\ &= \alpha + \beta\theta_t + \epsilon_{d,t} \end{aligned}$$

We use this idea to test whether the dealer knows the average market price for asset that he does not trade in a given month. Table 1.10 reports the result of this test for different outcome variables $y_{d,t}$ and instruments Z_t. The first two column test whether the dealer knows the average trading price of an asset in periods in which he does not trade. In all four of the combinations the p-value is zero, suggesting that the average price for the asset, $\theta_{t,a}$, or for assets from the same state $\theta_{t,s}$, don't belong to the dealer's information set when he does not trade. On the contrary, for periods in which the dealers did participate to trade the test cannot reject the null, confirming that dealer d acquire information through trade.

	$y_{d,t} = n^+_{d,t}$	$y_{d,t} = n^-_{d,t}$
$(1 - \mathbb{I}\{\text{trade in } t-1\})\,\theta_{t-1,a}$	0.00	0.00
$(1 - \mathbb{I}\{\text{trade in } t-1\})\,\theta_{t-1,s}$	0.00	0.00
$\mathbb{I}\{\text{trade in } t-1\}\,\theta_{t-1,a}$	0.02	0.75
$\mathbb{I}\{\text{trade in } t-1\}\,\theta_{t-1,s}$	0.10	0.15

Table 1.10: p-values for Hansen-Sargan test

1.C Definition of Θ

Below we outline the steps we follow to recover Θ. For every month t we compute the average trading price in trades between dealers and investors. Let $\hat{p}^+_t$ be the average price at which dealers buy the asset, and $\hat{p}^-_t$ be the average price at which dealers sell

the asset. According to our model, as the number of trades i increases,

$$\hat{p}_t^+ = \frac{1}{N}\sum p_{it}^+ \to \mathbb{E}\left(p_t^+|\theta_t\right) = \int p f\left(p|\theta_t,+\right) dp.$$
$$\hat{p}_t^- = \frac{1}{N}\sum p_{it}^- \to \mathbb{E}\left(p_t^-|\theta_t\right) = \int p f\left(p|\theta_t,+\right) dp.$$

To recover the underlying sequence for the parameter θ_t and its transition matrix h we fit a Normal Hidden Markov Model with three states to the sequence of market prices $\hat{p}_t = \frac{1}{2}\left(\hat{p}_t^+ + \hat{p}_t^-\right)$.

Concretely, we use an EM algorithm to select transition matrix $\hat{h}$ and initial probability $\hat{h}_0$ to maximize the expected log-likelihood

$$\mathbb{E}\left[l_T\left(p_{1:T}, \theta_{1:T}|h, h_0, \mu_\theta, \sigma_\theta\right)\right],$$

where

$$l_T\left(p_{1:T}, \theta_{1:T}|h, h_0, \mu_\theta, \sigma_\theta\right) = \log h_0\left(\theta_1\right) + \sum_{t=2}^{T} \log h\left(\theta_t|\theta_{t-1}\right) + \sum_{t=2}^{T} \log\left(\frac{1}{\sigma_{\theta_t}}\phi\left(\frac{\hat{p}_t - \mu_{\theta_t}}{\sigma_{\theta_t}}\right)\right).$$

Table 1.11 highlights the features of the recovered parameters across different classes of assets. The three columns of the table show, for each group of assets, the average purchase and selling prices by state, $\frac{1}{T}\sum p_t^{\cdot}\mathbb{I}\left\{\theta_t = \theta^k\right\}$. Finally the last two columns report the average number of changes for θ_t within an year and the average value of σ_{θ_t}. The state changes around three times within an year, this is a reasonable number considering that each group includes several assets. Furthermore the bid-ask spread is on average 5%, and it is larger for lower values assets (the correlation is -65%).

Group	Average Ask Price			Average Bid Price			Changes per year	σ_{θ_t}
	θ_{high}	θ_{mid}	θ_{low}	θ_{high}	θ_{mid}	θ_{low}		
I	1.02	1.038	1.053	0.997	1.014	1.03	4.06	0.015
II	0.701	0.738	0.776	0.707	0.714	0.727	1.029	0.038
III	1.007	1.017	1.026	0.989	0.995	1	3.383	0.015
IV	1.063	1.08	1.094	1.05	1.068	1.08	3.14	0.015
V	0.997	1.015	1.031	0.972	0.987	1.003	3.549	0.015
VI	1.044	1.049	1.054	1.029	1.035	1.04	3.071	0.014
VII	0.901	0.932	0.974	0.855	0.898	0.933	2.057	0.041
VIII	1.067	1.076	1.082	1.056	1.064	1.07	3.913	0.015
IX	0.949	0.98	0.997	0.911	0.943	0.957	1.662	0.029
X	0.948	0.964	0.983	0.914	0.931	0.945	2.903	0.017
XI	1.048	1.06	1.072	1.03	1.04	1.051	2.967	0.015
XII	0.801	0.82	0.837	0.778	0.796	0.813	3.429	0.021
XIII	1.029	1.03	1.036	1.012	1.013	1.017	3.086	0.014
XIV	0.975	0.993	1.018	0.943	0.963	0.983	3.464	0.013
XV	1.097	1.118	1.135	1.088	1.109	1.126	3.338	0.017

Table 1.11: Estimates of Θ

1.D Assets Classes

For the purpose of the estimation, we divide assets $j = \{1, \ldots, N_{\text{asset}}\}$ traded in the secondary market for municipal bonds into 15 groups, $a \in \{1, \ldots, 15\}$. Denote by $\bar{p}_{j,t}$ the average selling price asset j and denote by $\bar{p}_{a,t}$ the average selling price for assets belonging to class a in month t

$$\bar{p}_{a,t} = \frac{1}{N_a} \sum p_{i,t}.$$

Ideally the assignment of assets to classes should satisfy two conditions. First, for every class a past prices should have strong predictive power for the current price,

hence $\mathbb{C}ov\left(\bar{p}_{a,t}, \bar{p}_{a,t-k}\right)$ should be large. Moreover, knowing past trading price for class a', $\bar{p}_{a',t-k}$, should not help in predicting $\bar{p}_{a,t}$, conditional on the realization of the current public signal y_t^P.

To define classes that satisfy these conditions, we modify a standard k-means algorithm. Denote by $\mu^*(j) \in \{1, \ldots, 15\}$ the assignment of assets to classes. The algorithm follows these steps.

1. First we define a random assignment of assets $\mu^{(0)}$

2. To evaluate any assignment $\mu^{(m)}$ we first compute average prices within each class

$$\bar{p}_{a,t}^{(m)} = \frac{1}{N_a^{(m)}} \sum p_{\mu^{(m)}(j)=a,t},$$

   estimate the regression

$$\bar{p}_{a,t}^{(m)} = \rho_0 + \rho_1 \bar{p}_{a,t-1}^{(m)} + \rho_2 y_t^P + \epsilon, \tag{1.21}$$

3. Finally, we use 1.21 to update assignment $\mu^{(m)}$. In particular, $\mu^{(m+1)}(j)$ is defined as

$$\mu^{(m+1)}(j) = \arg\min_a \sum_t \left(\bar{p}_{j,t} - p_{a,t}^{(0)} - \hat{\rho}_0 - \hat{\rho}_1 \bar{p}_{a,t-1}^{(0)} - \hat{\rho}_2 y_t^P\right)^2.$$

Table 1.12 reports some characteristics of the different classes of assets. In particular, for the second column, we estimate separate regressions

$$\bar{p}_{a,t}^{(m)} = \rho_0 + \rho_{1,a'} \bar{p}_{a',t-1}^{(m)} + \rho_2 y_t^P + \epsilon,$$

and report the average p-value of the coefficients $(\rho_{1,a'})_{a' \neq a}$. Furthermore, for the third column, we estimate the regression

$$\bar{p}_{a,t}^{(m)} = \rho_0 + \rho_{1,a} \bar{p}_{a,t-1}^{(m)} + \rho_2 y_t^P + \epsilon,$$

and report the p-value associated to coefficient $\rho_{1,a}$.

Group	N. of assets	Total Volume of Trade ($\$10^9$)	$\rho_{1,a'}$	$\rho_{1,a}$
I	18,715	57.664	.377	.040
II	8,356	41.47	.539	.011
III	14,599	42.52	.586	0.001
IV	1,638	7.52	.456	0.215
V	14,876	87.99	.348	0
VI	14,371	61.86	.379	0
VII	17,145	59.44	.437	0
VIII	308	6.65	.327	0
IX	8,955	50.08	.383	0
X	1,275	11.87	.430	0
XI	19,266	65.72	.489	0.01
XII	1,952	6.08	.381	0
XIII	22,590	50.01	.541	0
XIV	9,516	90.39	.552	0.008
XV	7,498	87.65	.342	0

Table 1.12: Classes of Assets.

1.E Choice Probabilities

We assume that: (i) cost shocks $\left(\epsilon_{d,t}^{\mathbf{buy}}, \epsilon_{d,t}^{\mathbf{sell}}, \epsilon_{d,t}^{\varnothing}\right) \in \mathbb{R}^3$, are drawn from a double exponential distribution F_0, with standard deviation σ_0; (ii) $\epsilon_{d,t}^1 \in \mathbb{R}$ is drawn from a normal distribution F_1 with standard deviation σ_1; (iii) $\tilde{c}(n) = c_1 |n| + c_2 n^2$ is convex

in $|n|$; and (iv) experience transition matrix r can be rewritten as

$$r\left(e'|e,n,\tilde{e}\right)=\sum_{e'}r_{\text{d2d}}\left(e'|e'',\tilde{e}\right)r_{\text{inv}}\left(e''|e,n\right),$$

where r_{d2d} and r_{inv} describe, respectively, the change in experienced that can be attributed to inter-dealer trade trade with investors.

Denote by $V^{\text{sign}(n)}\left(\pi,x,e,\epsilon^1\right)$ the dealer's highest utility conditional on either buying or selling the asset:

$$\begin{aligned}&V^{\text{sign}(n)}\left(\pi,x,e,\epsilon^1\right)\\&=\max_{n\in\mathcal{N}(\text{sign}(n))}\left\{-\tilde{c}\left(n\right)-\epsilon_1\left|n\right|-\mathbb{E}\left(\sum_{i=1}^{|n|}p_{it}|\pi,\text{sign}\left(n\right)\right)+\mathbb{E}\left[W\left(\mathcal{L}_{inv}\left(\pi,\vec{p}_{n,t}\right),x'\left(n;x\right),r_{\text{inv}}\left(n,e\right)\right)\right]\right\},\end{aligned}$$

where

$$\mathcal{N}\left(\text{sign}\left(n\right)\right)=\begin{cases}\{1,\ldots,N\} & \text{sign}\left(n\right)=+1\\ \{-x,\ldots,-1\} & \text{sign}\left(n\right)=-1\end{cases},$$

The probability that dealer d chooses to trade with $n\neq 0$ investors can be written as

$$P\left(n_{d,t}=n|\pi,x,e\right)=\int_{\text{lb}(\pi,x,e,n)}^{\text{ub}(\pi,x,e,n)}\frac{\exp\left(\frac{V^{\text{sign}(n)}\left(\pi,x,e,\epsilon^1\right)-W(\pi,x,e)}{\sigma_0}\right)}{\exp\left(\frac{V^{\text{sign}(n)}(\pi,x,e,\epsilon^1)-W(\pi,x,e)}{\sigma_0}\right)+\exp\left(\frac{V^{\text{sign}(n)}(\pi,x,e,\epsilon^1)-W(\pi,x,e)}{\sigma_0}\right)+1}dF_1\left(\epsilon^1\right),$$

where $\text{ub}\left(\pi,x,e,n\right)$ and $\text{lb}\left(\pi,x,e,n\right)$ are optimal policy thresholds defined below. Consider $n>0$ and let $\Delta\left(\pi,x,e,n\right)$ denote the difference in the value function

between buying n and $n+1$ units:

$$\begin{aligned}\Delta(\pi, x, e, n) &= \tilde{c}(n+1) - \tilde{c}(n) - \mathbb{E}(p_{it}|\pi, +1), . \\ &\quad + \beta \mathbb{E} W\left(\mathcal{L}_{inv}(\pi, \vec{p}_{n+1}), x'(n+1; x), r_{\text{inv}}(n+1, e)\right) \\ &\quad - \beta \mathbb{E} W\left(\mathcal{L}_{inv}(\pi, \vec{p}_{n}), x'(n; x), r_{\text{inv}}(n+1, e)\right).\end{aligned}$$

Then

$$\sigma_1 \text{lb}(\pi, x, e, n) = \begin{cases} \Delta(n) & n = 1, \ldots, N-1 \\ -\infty & n = N \end{cases},$$

while

$$\sigma_1 \text{ub}(\pi, x, e, n) = \begin{cases} \Delta(n-1) & n = 2, \ldots, N \\ \infty & n = 1 \end{cases}.$$

In the same fashion, consider $n < 0$ and let $\Delta(\pi, x, e, n)$ denote the difference in the value function between selling n and $n+1$ units:

$$\begin{aligned}\Delta(\pi, x, e, n) &= \tilde{c}(n-1) - \tilde{c}(n) + \mathbb{E}(p_{it}|\pi, -1), \\ &\quad + \beta \mathbb{E} W\left(\mathcal{L}_{inv}(\pi, \vec{p}_{n-1}), x'(n-1; x), r_{\text{inv}}(n-1, e)\right) \\ &\quad - \beta \mathbb{E} W\left(\mathcal{L}_{inv}(\pi, \vec{p}_{n}), x'(n; x), r_{\text{inv}}(n-1, e)\right).\end{aligned}$$

Then

$$\sigma_1 \text{lb}\left(\pi, x, e, n\right) = \begin{cases} \Delta\left(n\right) & n = -1, \ldots, -x-1 \\ -\infty & n = -x \end{cases},$$

while

$$\sigma_1 \text{ub}\left(\pi, x, e, n\right) = \begin{cases} \Delta\left(n-1\right) & n = -2, \ldots, -x \\ \infty & n = -1 \end{cases}.$$

1.F Estimation Algorithm

1. Guess an initial set of parameters τ.

2. Solve for the dealers' value functions. This requires setting an initial value $V_{(0)}$. Then, at each iteration m and until convergence

   (a) Using the observed distribution of beliefs, experience, and inventory recovered in Section 1.5.2, update $W^{\text{buy}}_{(m)}$according to

$$\begin{aligned} &W^{\text{buy}}_{(m)}\left(\pi, x, e\right) \\ &= \mathbb{P}\left(\tilde{e} = 0 | \pi, x, e\right) \beta \mathbb{E}\left[V_{(m)}\left(\pi, x, g_2\left(0; e\right)\right)\right] - \hat{\mathbb{E}}\left(\tilde{q} | \pi, x, e\right) + \\ &+ \sum_{\tilde{e}=1}^{\bar{E}} \mathbb{P}\left(e, \tilde{e}\right) \int \beta \mathbb{E}\left[V_{(m)}\left(\mathcal{L}_{\text{buy}}\left(\tilde{y}, \tilde{q}, e, \tilde{e}\right), x', g_2\left(\tilde{e}; e\right)\right) | e, \tilde{e}, \tilde{q}\right] \hat{f}_q\left(\tilde{q} | \tilde{e}, \pi, x, e\right) d\tilde{q}. \end{aligned}$$

   Note that $\hat{\mathbb{E}}\left(\tilde{q} | \pi, x, e\right)$ is the actual average price at which a dealer in state (π, x, e) buys the asset. Similarly, we use the observed distribution of accepted offers $\hat{f}_q$ and observed matching probabilities $\mathbb{P}\left(e, \tilde{e}\right)$.

(b) Update $\tilde{W}^{\text{sell}}_{(m)}$:

$$\begin{aligned}
&\tilde{W}^{\text{sell}}_{(m)}(\pi, x, e, \tilde{e}) \\
&= \max_{q} \mathbb{P}_{(m)}(q \text{ is accepted}|e, \tilde{e})\, \beta \mathbb{E}\left[V_{(m)}\left(\mathcal{L}_{\text{sell}}(\tilde{y}, q, e, \tilde{e}), \tilde{x}', g_2(\tilde{e}; e)\,|e, \tilde{e}, q\right)\right] \\
&+ \mathbb{P}_{(m)}(q \text{ is rejected}|e, \tilde{e})\, \beta \mathbb{E}\left[V_{(m)}(\pi, x, g_2(0; e))\right],
\end{aligned}$$

(c) Update $W^{\text{sell}}_{(m)}$:

$$\begin{aligned}
&W^{\text{sell}}_{(m)}(\pi, x, e) \\
&= \sigma_v \log\left(\sum_{\tilde{e}=1}^{\bar{E}} \exp\left(\frac{-c(\tilde{e}) + \tilde{W}^{\text{sell}}_{(m)}(\pi, x, e, \tilde{e})}{\sigma_v}\right) + \exp\left(\frac{\beta \mathbb{E}\left[V_{(m)}(\pi, x, g_2(0; e))\right]}{\sigma_v}\right)\right) + \sigma_v \gamma.
\end{aligned}$$

(d) Update $W_{(m)}$:

$$W_{(m)} = \alpha W^{\text{sell}}_{(m)} + (1 - \alpha) W^{\text{buy}}_{(m)}.$$

(e) Update $V_{(m+1)}$:

$$\begin{aligned}
&V_{(m+1)}(\pi, x, e) \\
&= -\kappa(x) + \mathbb{E}\left(\max_{n \in \{-x, \ldots, 0, 1, \ldots, N\}} \left\{\begin{array}{c} -c(n, \epsilon) - \text{sign}(n)\, \mathbb{E}\left(\sum_{i=1}^{|n|} p_{it} | \pi, \text{sign}(n)\right) \\ + \mathbb{E}\left[W\left(\mathcal{L}_{inv}(\pi, (p_i)_{i=1}^{n}, \text{sign}(n)), x'(n; x), g_2(|n|; e)\right)\right] \end{array}\right\}\right).
\end{aligned}$$

3. Compute optimal choice probabilities according to (1.7) and (1.11), and simulate choices $\left\{\mathbb{I}^{(m)}_{d.t,\tilde{e}}, \mathbb{I}^{(m)}_{d.t,\tilde{e}}\right\}_{m=1}^{M}$.

4. For each simulated sample, find the auxiliary parameters that maximize $\mathcal{L}\left(\tilde{\mathbb{I}}; \alpha, \beta, z\right)$.

1.G Steady-State conditions

Let f^*_{inv} denote the distribution of dealers' private history after trade with investors implied by f^*. Then, f^*_{inv} satisfies

$$\begin{aligned}
&f^*_{\text{inv}}\left(\pi_d \in A_\pi, x_d \in A_x, e_d \in A_e|\theta\right) \\
&= \sum_{n=-N}^{N} \int \left(\int \mathbb{I}_{A(\vec{p}_n)} f\left(\vec{p}_n|\theta, \text{sign}\left(n\right)\right) d\vec{p}_n \right) P\left(n|\pi_d, x_d, e_d\right) f^*\left(\pi_d, x_d, e_d\right) d\left(\pi_d, x_d, e_d\right),
\end{aligned}$$

where

$$A\left(\vec{p}_n\right) = \left\{\mathcal{L}_{inv}\left(\pi_d; \vec{p}_n, \text{sign}\left(n\right)\right) \in A_\pi, x_d - n \in A_x, r_{\text{inv}}\left(e_d, n\right) \in A_e\right\},$$

the probability that dealer d chooses n, $P\left(n|\pi_d, x_d, e_d\right)$, is defined in 1.7, and $\mathcal{L}_{inv}$ is defined in 1.3.

Next, let $f^*_q\left(q \in A_q|e, \tilde{e}, \theta\right)$ be the equilibrium distribution of offers directed from a seller of type e to a buyer of type $\tilde{e}$, in state θ. Then

$$\begin{aligned}
&f^*_q\left(q \in A_q|e, \tilde{e}, \theta\right) \\
&= \int \int \mathbb{I}\left\{q\left(\pi_d, x_d, e_d, \tilde{e}\right) \in A_q\right\} f^*_{\text{inv}}\left(\pi_d, x_d, e_d|\theta, P_e\left(\pi_d, x_d, e_d, \xi\right) = \tilde{e}\right) d\left(\pi_d, x_d, e_d\right) d\xi,
\end{aligned}$$

where $q\left(\cdot\right)$ achieves the maximum in (1.9), and $P_e\left(\cdot\right)$ solves (1.10).

Finally, define $f^*_r\left(r|e, \tilde{e}, \theta, q\right)$ to be the equilibrium probability that an offer q receives reply $r \in \{0, 1\}$, conditional on the seller having experience e, the buyer having experience $\tilde{e}$, and the state being θ:

$$f^*_r\left(\bar{r}|e, \tilde{e}, \theta, q\right) = \int \int \mathbb{I}\left\{r\left(\pi_d, x_d, e, \tilde{e}, q\right) = r\right\} f^*_{\text{inv}}\left(\pi_d, x_d|\theta, e_d = e\right) d\left(\pi_d, x_d\right),$$

where $r\left(\cdot\right)$ achieves the maximum in (1.8).

Then the steady state distribution f^* must satisfy

$$f^*\left(\pi_d \in A_\pi, x_d \in A_x, e_d \in A_e|\theta\right) = \alpha f^*_{\text{sell}}\left(\pi_d \in A_\pi, x_d \in A_x, e_d \in A_e|\theta\right) + (1-\alpha) f^*_{\text{buy}}\left(\pi_d \in A_\pi, x_d \in A_x, e_d \in A_e|\theta\right),$$

where $f^*_{\text{sell}}(\cdot|\theta)$ and $f^*_{\text{buy}}(\cdot|\theta)$ are the distributions of the state among potential buyers and sellers, after inter-dealer trade. In particular

$$\begin{aligned} &f^*_{\text{sell}}\left(\pi_d \in A_\pi, x_d \in A_x, e_d \in A_e|\theta\right) \\ &= \sum_{\tilde{e}=1}^{\bar{E}} \int\int \mathbb{I}_{A_{\text{sell}}(\pi_d, x_d, e_d, \tilde{e}, r, q)} f^*_r\left(r|e_d, \tilde{e}, \theta, q\right) f^*_q\left(q|e_d, \tilde{e}, \theta\right) P\left(\tilde{e}|\pi_d, x_d, e_d\right) f^*_{inv}\left(\pi_d, x_d, e_d\right) d\left(\pi_d, x_d, e_d\right) \\ &+ \int \mathbb{I}_{A_{\text{no trade}}} P\left(0|\pi_d, x_d, e_d\right) f^*_{inv}\left(\pi_d, x_d, e_d|\theta\right) d\left(\pi_d, x_d, e_d\right), \end{aligned}$$

where

$$A_{\text{no trade}} = \left\{\pi_d \in A_\pi, x_d \in A_x, e_d \in A_e\right\},$$

$$A_{\text{sell}}\left(\pi_d, x_d, e_d, \tilde{e}, r, q\right) = \left\{\mathcal{L}_{sell}\left(\pi_d, r, q, e_d, \tilde{e}\right) \in A_\pi, x_d - r \in A_x, r_{\text{d2d}}\left(e_d, \tilde{e}\right) \in A_e\right\},$$

and $f^*_{\text{buy}}(\cdot|\theta)$ can be defined similarly.

1.H Proofs from Section 1.5

First rewrite

$$\begin{aligned} c_t(\delta) &= \sum_{k\geq 1} \delta^k I_{t,t-k} \\ &= \delta\alpha g_t n_{t-1} + \delta^2\alpha^2 g_t g_{t-1} n_{t-2} \dots \\ &= \delta\alpha g_t \left(n_{t-1} + \sum_{k\geq 2} \delta^{k-1}\alpha^{k-1} g_{t-1}\dots n_{t-k} \right). \end{aligned}$$

Next, note that

$$\begin{aligned} r_t &= n_{t-1} + \delta\alpha g_{t-1}\left(n_{t-2} + \delta\alpha g_{t-2} r_{t-2}\right) \\ &= n_{t-1} + \delta\alpha g_{t-1} n_{t-2} \\ &+ \delta^2\alpha^2 g_{t-1} g_{t-2}\left(n_{t-3} + \delta\alpha g_{t-3} r_{t-3}\right) \\ &= \cdots \\ &= n_{t-1} + \sum_{k\geq 2} (\delta\alpha)^{k-1} g_{t-1}\cdots g_{t-k} n_{t-k}, \end{aligned}$$

then

$$e_t - \delta e_{t-1} = n_t + \delta\alpha g_t r_t.$$

Simply note that the sequence $(r_t)_{t\geq 0}$, satisfies

$$\begin{aligned} r_t &= n_{t-1} + \sum_{k\geq 2} (\delta\alpha)^{k-1} g_{t-1}\cdots g_{t-k} n_{t-k} \\ &< N + \sum_{k=1}^{t} (\delta\alpha)^k D^k N^{k-1} \end{aligned}$$

Therefore,

$$\begin{aligned}
e_t &= \sum_{k\geq 1} \delta^k \left(N + \sum_{h=1}^{k} (\delta\alpha)^h D^h N^{h-1} \right) \\
&= \sum_{k\geq 1} \delta^k N + \sum_{k\geq 1} \delta^k \sum_{h=1}^{k} (\delta\alpha)^h D^h N^{h-1} \\
&< \frac{N}{1-\delta} + \alpha D \sum_{k\geq 1} \delta^{k+1} \frac{1-(\delta\alpha)^{k+1} D^{k+1} N^{k+1}}{1-(\delta\alpha) DN} \\
&< \frac{N}{1-\delta} + \frac{\alpha D}{1-(\delta\alpha) DN} \left(\frac{1}{1-\delta} - \delta^4 \alpha^2 (DN)^2 \sum_{k\geq 0} \delta^{2k} \alpha^k D^k N^k \right).
\end{aligned}$$

This, in turn, is bounded if

$$\delta^2 \alpha < \frac{1}{DN}.$$

Focus on the situation of seller d who trades with buyer $\tilde{d}$ at price $q_{d,t}$. Denote by $e_{d,t}$ and $e_{\tilde{d},t}$, respectively, dealer d and dealer $\tilde{d}$'s experience levels. Furthermore $\hat{\pi}_{d,t}$ and $\hat{\pi}_{\tilde{d},t}$ denote dealers d and $\tilde{d}$ beliefs at the moment of the trade (that is, after inter-dealer trade). Dealer d observes post-trade signal $y_{\tilde{d},t} = \pi_{\tilde{d},t}$, as well as dealer $\tilde{d}$'s decision about whether to accept the offer, $r_{\tilde{d},t} = r\left(\pi_{\tilde{d},t}, x_{\tilde{d},t}, e_{\tilde{d},t}, q_{d,t}, e_{d,t}\right)$. Dealer d's updated belief satisfies

$$\hat{\hat{\pi}}_{d,t}\left(\theta_t = \theta^k | \pi_{\tilde{d},t}, e_{\tilde{d},t}, q_{d,t}, e_{d,t}\right) = \frac{f^*\left(\pi_{\tilde{d},t}, r_{\tilde{d},t} | e_{\tilde{d},t}, q_{d,t}, e_{d,t}, \theta^k\right) \hat{\pi}_{d,t}\left(\theta^k\right)}{\sum_\theta f^*\left(\pi_{\tilde{d},t}, r_{\tilde{d},t} | e_{\tilde{d},t}, q_{d,t}, e_{d,t}, \theta\right) \hat{\pi}_{d,t}\left(\theta\right)}.$$

We can rewrite this as

$$\begin{aligned}
& f^*\left(\pi_{\tilde{d},t}, r\left(\pi_{\tilde{d},t}, x_{\tilde{d},t}, e_{\tilde{d},t}, q_{d,t}, e_{d,t}\right) | e_{\tilde{d},t}, q_{d,t}, e_{d,t}, \theta^k\right) \\
= & \int_{A_r\left(\pi_{\tilde{d},t}, e_{\tilde{d},t}, q_{d,t}, e_{d,t}\right)} f^*\left(\pi_{\tilde{d},t}, x_{\tilde{d},t} | e_{\tilde{d},t}, \theta^k\right) dx_{\tilde{d},t},
\end{aligned}$$

where

$$A_r\left(\pi_{\tilde{d},t}, e_{\tilde{d},t}, q_{d,t}, e_{d,t}\right) = \left\{x_{\tilde{d},t} : r\left(\pi_{\tilde{d},t}, x_{\tilde{d},t}, e_{\tilde{d},t}, q_{d,t}, e_{d,t}\right) = r\right\}$$

However information about common shock θ_t contained in inventory $x_{\tilde{d},t}$ is already incorporated in $\pi_{\tilde{d},t}$, since dealer $\tilde{d}$ knows $x_{\tilde{d},t}$. Therefore

$$\begin{aligned} & f^*\left(\pi_{\tilde{d},t}, x_{\tilde{d},t} | e_{\tilde{d},t}, \theta^k\right) \\ =& f^*\left(x_{\tilde{d},t} | e_{\tilde{d},t}, \pi_{\tilde{d},t}\right) f^*\left(\pi_{\tilde{d},t} | e_{\tilde{d},t}, \theta^k\right), \end{aligned}$$

and

$$\begin{aligned} & f^*\left(\pi_{\tilde{d},t}, r\left(\pi_{\tilde{d},t}, x_{\tilde{d},t}, e_{\tilde{d},t}, q_{d,t}, e_{d,t}\right) | e_{\tilde{d},t}, q_{d,t}, e_{d,t}, \theta^k\right) \\ =& f^*\left(\pi_{\tilde{d},t} | e_{\tilde{d},t}, \theta^k\right) \int_{A_r\left(\pi_{\tilde{d},t}, e_{\tilde{d},t}, q_{d,t}, e_{d,t}\right)} f^*\left(x_{\tilde{d},t} | e_{\tilde{d},t}, \pi_{\tilde{d},t}\right) dx_{\tilde{d},t}. \end{aligned}$$

Since the last term doesn't depend on common shock θ_t, we can rewrite

$$\hat{\hat{\pi}}_{d,t}\left(\theta_t = \theta^k | \pi_{\tilde{d},t}, e_{\tilde{d},t}, q_{d,t}, e_{d,t}\right) = \frac{f^*\left(\pi_{\tilde{d},t} | e_{\tilde{d},t}, \theta^k\right) dx_{\tilde{d},t} \hat{\pi}_{d,t}\left(\theta^k\right)}{\sum_\theta f^*\left(\pi_{\tilde{d},t} | e_{\tilde{d},t}, \theta\right) \hat{\pi}_{d,t}\left(\theta\right)}.$$

Chapter 2

Geography, Search Frictions and Endogenous Trade Costs[1]

2.1 Introduction

More than 80% of international trade is carried by the global shipping industry.[2] To export, an exporter has to find an available vessel and contract for a voyage and a price.[3] In turn, the ship is optimally choosing its travels in search of cargo, thinking about its future options. This spatial equilibrium between exporters and ships determines the trade costs (i.e. transportation prices) that different countries face, as well as the trade flows between different regions of the world. What is the role of geography (i.e. country locations, natural inheritance of goods) in determining

[1]The paper on which this chapter is based is co-authored with Myrto Kalouptsidi and Theodore Papageorgiou. A version of this paper was submitted for publication to Econometrica and the journal editor has invited us to revise the paper and resubmit. The paper has been presented at AEA, NBER, Cowles confererence, Barcelona Summer Forum, SED, Darmuth Winter IO conference, and at "Implementation of Structural Dynamic Models: Methodology and Applications".

[2]Source: International Chamber of Shipping. Seaborne trade accounts for about 70% of trade in terms of value.

[3]Different segments of shipping function differently. In this paper we focus on bulk shipping, where exporters of bulk commodities fill up an entire vessel and hire it for a single trip in a spot market (much like a taxi or a rental car); see Section 2.2.2.

trade costs and flows? Is the matching process between exporters and ships efficient? How does ship behavior affect the behavior of exporters and the resulting trade flows?

In this paper, we use detailed micro-data on vessel movements, as well as shipping contracts between exporters and ships to shed new light on world trade costs and trade flows. The data both reveal novel facts about shipping patterns, and motivate us to build a framework modeling the behavior of exporters and transportation agents (ships). Our framework has two novel features. First, trade costs are endogenous and determined jointly with trade flows. As such, trade costs depend on the entire network of countries, rather than just the bilateral (distance between) trading partners. Endogenous trade costs provide a novel link to understand trade patterns. Second, search frictions between exporters and ships can limit trade. We use the data to estimate our model and recover flexibly the matching process between ships and exporters. Finally, we use our framework to tackle a number of questions of interest: How does an improvement in shipping efficiency affect world trade flows? How do shocks propagate through the world; for instance, how would a Chinese slow-down trickle through the network of countries, or how would the opening of the Northwest Passage affect trade costs and trade flows? What is the loss due to search frictions between exporters and ships?

We focus on dry bulk ships, which carry mostly commodities (grain, ore, coal, etc.) and whereby an exporter hires the entire vessel for a specific voyage. We begin by using our data to uncover some novel facts about (i) world trade flows; (ii) trade costs; (iii) the matching process between ships and exporters. First, satellite data of ships' movements reveal that most countries are either net importers or net exporters and that, related to this, at any point in time a staggering 45% of ships are traveling without cargo (ballast). This natural trade imbalance, often overlooked in the trade literature, is a key driver of trade costs. Second, shipping prices are largely asymmetric and depend on the destination's trade imbalance: all else equal,

the prospect of having to ballast after the destination port leads to higher prices. For instance, shipping from Australia to China is 30% more expensive than the reverse; as China mostly imports raw materials, ships arriving there have limited opportunities to reload. Third, we uncover evidence suggestive of search frictions between ships and exporters: at any given time, in most countries there are simultaneous arrivals of empty ships that load and departures of empty ships, even though ships are homogeneous. Moreover, the law of one price does not hold: shipping prices exhibit substantial dispersion within a time-origin-destination triplet.

We build a dynamic spatial search model of the global shipping industry, in the spirit of the search and matching literature, that captures the observed empirical patterns and explores the importance of endogenous trade costs. The model features three key ingredients: (i) geography; (ii) search frictions; (iii) forward-looking ships and exporters that optimally choose their travels and exporting destinations respectively. Geography enters the model through different trip durations across different locations. In addition to natural geography, locations differ in their economic geography, namely their natural inheritance in commodities of different value. Search frictions between exporters and ships prevent the matching of all possible pairs. Ships are homogeneous and forward looking: when matched with an exporter, they negotiate the price taking into account matching opportunities and ballast options at the destination. When unmatched, ships decide whether to wait at their location or ballast someplace else, taking into account their expected discounted stream of profits at each location.

We estimate our model using the collected data. There are two sets of core model primitives: (i) the matching function between ships and exporters, as well as the global distribution of searching exporters; and (ii) ship sailing and port waiting costs, as well as exporter valuations and costs. Using data on the number of ships and

matches, as well as the weather, we obtain the former; using data on shipping prices, ship ballast choices and exporter destination choices, we estimate the latter.

In particular, we adopt a novel approach to flexibly recover both the matching function (which gives the number of matches as a function of the number of agents searching on each side of the market), as well as the searching exporters. A sizable literature has estimated matching functions in different contexts (e.g. labor markets, taxicabs).[4] Here, we observe ships and matches, but not searching exporters. Our approach draws from the literature on nonparametric identification (Matzkin (2003)) and, to our knowledge, we are the first to apply it to matching function estimation. It relies on the joint density of matches and ships, as well as sea weather as an instrument that exogenously changes arriving ships. We make two contributions. First, unlike the existing literature, we do not take a stance on the presence of search frictions. When one side of the market (in this case exporters) is unobserved or mismeasured, it is difficult to discern whether search frictions are present. Second, we avoid parametric restrictions on the matching function; this is important, since in frictional markets, the shape of the underlying matching function is directly linked to welfare implications. To provide some intuition, consider the following test for search frictions: weather shocks exogenously shift ship arrivals at port; in regions with more ships than exporters, this should not affect matches unless there are search frictions. We show that here, matches are indeed affected by weather shocks, which both suggests that search frictions are present and allows us to recover the curvature of the matching function.

The remaining primitives are obtained from ship and exporter choices, and prices. In particular, we first recover ship sailing and port costs via Maximum Likelihood, formed by the optimal ballast choice probabilities. As ships are forward looking

[4]For instance, in labor markets, data on unemployed workers, vacancies and matches delivers the underlying matching function. In the market for taxi rides, one observes taxis and their rides, but not hailing passengers; in recent work, Frechette et al. (2016) and Buchholz (2016) have used a parametric assumption on the matching function, to recover the passengers.

this is a dynamic discrete choice problem (similar to Rust (1987)). Then, we obtain exporter valuations directly from prices: once ship primitives are known, we employ the surplus sharing condition derived from Nash bargaining to back out each valuation corresponding to each individual contract price. Finally, we use trade flows (loaded trips) to recover exporter costs by destination.

Our estimated model provides a unique framework to study how trade costs and flows are jointly determined. The novel feature here is that trade costs are endogenous and we perform a number of counterfactual exercises to illustrate why this matters. More specifically, our counterfactuals showcase three mechanisms present in our economy: First, changes in the model's primitives also affect ships' outside opportunities and this has an indirect effect on prices and exports. Second, the change in a country's trade costs depends heavily on its network of trading partners and its geographical proximity to large exporters and importers. Third, reductions in impediments to trade, such as improvements in transportation efficiency or a reduction in search frictions, benefit disproportionally exporters who are (i) large; and (ii) export high value commodities, as ships are more likely to reallocate there. As a result this leads to "polarization", whereby differences in export volumes of different countries widen.

We first consider a decline in the cost of sailing (improvement in shipping efficiency). This naturally increases the value of a match between a ship and an exporter and thus pushes down the price while increasing exports. However, the decline in the sailing cost improves a ship's bargaining position, as it makes ballasting less costly. This latter effect dampens the original increase in exporting, and it pushes up the exports of net exporters disproportionately more than the exports of net importers.

We illustrate the importance of trade networks and market conditions in neighboring countries, by considering a slow-down in China. In such an event, the reallocation of ships over space can amplify the effect of the slow-down in neighboring regions.

More specifically, besides the direct effect to countries whose exports relied heavily on the Chinese economy, our model points out that there is a secondary effect driven by the reduced supply of ships in that region of the world: the many ships that ended up in China are no longer around. This impacts negatively both China's own exports (import-export complementarity), but also neighboring countries' toward which these ships would ballast; in contrast, distant countries may benefit from ships reallocating there.

We also consider the opening of the Northwest Passage. Melting the arctic ice would reduce the travel costs between Northeast America and the Far East; although the former experiences an increase in exporting, the latter suffers a decline because of the higher outside option of ships ending a trip there. Moreover, although the shock is local, it has global effects: other countries' exports are also lower due to ships' higher outside options. Exporters close to Northeast America (e.g. Brazil) are disproportionately hurt, as ships that used to ballast there now ballast to Northeast America; in contrast, other big exporters, such as Australia, are shielded by their closeness to the biggest importers (China, India).

Finally, we quantify the trade lost due to search frictions. We demonstrate that exporting universally goes up considerably. In addition, trade shifts towards countries with bigger exporters, as differences in frictions across regions are no longer relevant and exporter size becomes a more important determinant of trade.

Related Literature We relate equally to three broad strands of literature: (i) trade and geography; (ii) search and matching; (iii) industry dynamics.

First, our paper endogenizes trade costs and so it naturally relates to the large literature in international trade studying the importance of trade costs in explaining trade flows between countries (e.g. Anderson and Van Wincoop (2003), Eaton and Kortum (2002)). In much of this literature, trade costs are treated as "residuals"

that explain the gap between actual and predicted bilateral trade flows conditional on variables such as size, distance, common border/language, etc. Here, we consider what happens to trade flows when transport prices (an important component of trade costs, at least as large or larger than tariffs; Hummels (2007)) are determined in equilibrium, jointly with trade flows. In addition, we document important features of trade costs, often not taken into account; for instance, trade costs are asymmetric and depend on the trip's origin and destination, as well as the entire country network. Waugh (2010) has argued that asymmetric trade costs are necessary to explain some empirical regularities regarding trade flows across rich and poor countries.

We also contribute to a literature that has considered the role and features of the (container) shipping industry; e.g. Hummels and Skiba (2004) explore the relationship between product prices at different destinations and shipping costs; Hummels et al. (2009) explore market power in container shipping; Ishikawa and Tarui (2015) theoretically investigate the impact of "backhaul" and its interaction with industrial policy; Asturias (2016) explores the impact of the number of shipping firms on transport prices and trade; Wong (2017) incorporates container shipping prices featuring a "round-trip" effect in a trade model. Finally, recent work has explored the matching of importers and exporters under frictions (Eaton et al. (2016)).

Our paper is also related to both old and new work on the role of geography in international trade (e.g. Krugman (1991), Head and Mayer (2004), Allen and Arkolakis (2014)), as well as the impact of transportation infrastructure and networks (e.g. Donaldson (2012), Allen and Arkolakis (2016), Donaldson and Hornbeck (2016), Fajgelbaum and Schaal (2017)). We extend this literature by exploring how the location and neighborhood of each country interact with the functioning of transportation agents and thus shapes up trade costs and flows. We illustrate that both a country's location (i.e. distances from all other countries) and its natural inheritance are key features of the equilibrium. This natural asymmetry contributes to trade imbalances

which are often overlooked in the literature (one exception is Reyes-Heroles (2016)); as we argue however, these imbalances are crucial in determining a country's exports and trade costs. One feature of the trade literature that our paper is missing is that we do not determine product and input prices in equilibrium. The latter may be reasonable as we focus only on commodities and so wages and capital prices may be taken as exogenous. The former would require additional data and is an interesting avenue for future research.

Second, our paper relates to the search and matching literature (see Rogerson et al. (2005) for a survey). On one hand, our model is essentially a search model in the spirit of Mortensen and Pissarides (1994) where firms and workers (randomly) meet subject to search frictions and Nash bargain over a wage. An important addition in our case is the spatial nature of our setup: there are several interconnected markets at which agents (ships) can search. Lagos (2000), Lagos (2003) and Buchholz (2016) have also adopted similar spatial search models for taxi cabs; an important difference here is that prices are set in equilibrium, while in the taxi market prices are exogenously set by regulation. In our setup this is important, since by endogenizing the trade costs we can consider how they change trade flows in the different counterfactuals. Moreover, as mentioned above, our paper also contributes to the literature on matching function estimation (see Petrongolo and Pissarides (2001) for a survey).

Third, we relate to the literature on industry dynamics (e.g. Hopenhayn (1992), Ericson and Pakes (1995)). Consistent with this research agenda, we study the long-run industry equilibrium properties, in our case the spatial distribution of ships and exporters. Moreover, our empirical methodology borrows from the literature on the estimation of dynamic setups (e.g. Rust (1987), Aguirregabiria and Mira (2007), Bajari et al. (2007), Pakes et al. (2007)) in matching conditional choice probabilities that involve value functions (applications include Ryan (2012) and Collard-Wexler (2013)). Buchholz (2016) and Frechette et al. (2016) also explore dynamic decisions

in the context of taxi cabs' search and shift choices respectively. Finally, Kalouptsidi (2014) has also looked at the shipping industry, albeit at the entry decisions of shipowners and the resulting investment cycles in new ships, while Kalouptsidi (2017) focuses on industrial policy in the Chinese shipbuilding industry.

The rest of the paper is structured as follows. Section 2.2 provides a description of the industry and the data used. Section 1.3 presents a number of facts on trade flows, transportation prices and search frictions. Section 1.4 describes the model. Section 2.5 lays out our empirical strategy, while Section 1.6 presents the estimation results. Section 2.7 discusses the counterfactuals, while Section 1.9 concludes. The Appendix contains additional tables and figures, proofs to our propositions, as well as further data and estimation details.

2.2 Industry and Data Description

2.2.1 Trade in Dry Bulk Commodities

Dry bulk shipping involves vessels designed to carry a homogeneous unpacked dry cargo, for individual shippers on non-scheduled routes. The entire cargo carried belongs to one owner (the exporter). Bulk carriers operate like taxi cabs: a specific cargo is transported individually in a specific ship, for a trip between a single origin and a single destination. Dry bulk shipping involves mostly commodities, such as iron ore, steel, coal, bauxite, phosphates, but also grain, sugar, chemicals, lumber and wood chips; it accounts for about half of total seaborne trade in tons (UNCTAD) and 45% of the total world fleet, which includes also containerships and oil tankers.[5][6]

[5]It is worth noting that bulk ships are very different from containerships, which operate like buses: containerships carry cargos (mostly manufactured goods) from many different cargo owners in container boxes, along fixed itineraries/routes according to a timetable. It is technologically impossible to substitute bulk with container shipping.

[6]It is not straightforward to obtain information on the share of world trade value carried by bulkers. However, mining, agricultural products, chemicals and iron/steel jointly account for about 30% of world trade value (WTO (2015)).

There are four different categories of dry bulk carriers based on size: Handysize (10,000–40,000 DWT), Handymax (40,000–60,000 DWT), Panamax (60,000–100,000 DWT) and Capesize (larger than 100,000 DWT). The industry is unconcentrated, consisting of a large number of small shipowning firms (see Kalouptsidi (2014)): the maximum fleet share is around 4%, while the firm size distribution features a large tail of small shipowners. Moreover, shipping services are largely perceived as homogeneous. In his lifetime, a shipowner will contract with hundreds of different exporters, carry a multitude of different products and visit numerous countries; we discuss this further in Section 1.3.

Trips are realized through individual contracts: shipowners have vessels for hire, exporters have cargo to transport and brokers put the deal together. Ships carry at most one freight at a time: the exporter fills up the hired ship with his cargo. In this paper, we focus on spot contracts and in particular the so called "trip-charters", in which the shipowner is paid in a per day rate.[7] The exporter who hires the ship is responsible for the trip costs (e.g. fueling), while the shipowner incurs the remaining ship costs (e.g. crew, maintenance, repairs).

2.2.2 Data

We combine a number of different datasets. First, we employ a dataset of dry bulk shipping contracts, from 2001 to 2016, collected by Clarksons Research, a major global shipbroking firm. An observation is a transaction between a shipowner and a charterer, for transportation of a specific cargo, on specific dates, from a specific origin to a specific destination. We observe the name of the vessel, the identity of the charterer who hires the ship, the contract signing date, the agreed loading

[7] Trip-charters are the most common type of contract. Long-term contracts ("time-charters"), however, do exist: on average, about 10% of contracts signed are long-term. Interestingly, though, ships in long-term contracts, are often "relet" in a series of spot contracts, suggesting that arbitrage is possible.

and unloading dates, the agreed upon trip price in dollars per day, as well as some information on the origin and destination.

Second, we use satellite AIS (Automatic Identification System) data from exactEarth Ltd (henceforth EE) for the ships in the Clarksons dataset between 2009 and 2016. AIS transceivers automatically broadcast information, such as the ships' positions (longitude and latitude), speed, and level of draft (i.e. the vertical distance between the waterline and the bottom of the ship's hull), at regular intervals of at most six minutes. The level of draft allows us to determine whether a ship is loaded or not at any point in time.

We also use the ERA-Interim archive, from the European Centre for Medium-Range Weather Forecasts (CMWF), to collect global data on daily sea weather. This allows us to construct weekly data on the wind speed (in each direction) on a 0.75 grid across all oceans. Finally, we employ several time series from Clarksons on e.g. the total fleet and fuel prices, as well as country-level imports/exports, production and commodity prices from numerous sources (e.g. UNCTAD, FAO, IEA, Comtrade).

	Mean	St. Dev	Median	Min	Max
Contract price per day (10^5 US dollars)	0.105	0.054	0.095	0.01	0.7
Contract trip price (10^5 US dollars)	1.417	9	1.17	0.07	8.315
Contracts per ship	2.108	1.445	2	1	12
Loaded trip duration (weeks)	2.50	2.06	2.02	0.14	10.41
Empty trip (ballast) duration (weeks)	1.74	1.63	1.29	0.09	8.98
Days between contract signing and loading date	6.11	6.686	4	0	28
Prob of ship staying at port conditional on being unmatched	0.77	0.12	0.76	0.59	0.95

Table 2.1: Summary statistics.

Summary statistics Our final dataset involves 5,410 ships between 2012 and 2016.[8] We end up with 7,652 shipping contracts, for which we know the price,

[8]We only use contracts during the sample period of the satellite data. Moreover, we drop the first two years (2009-2011) in the matching function estimation, as satellites are still launched at that time and the geographic coverage is more limited.

as well as the exact origin and destination. As shown in Table 2.1, the average price is 10,000 dollars per day (or 140,000 dollars if we take the trip duration into account), with substantial variation (the standard deviation is 5,000 dollars per day). We have 233,580 ship-week observations at which the ship decides to either ballast someplace or stay at its current location. Loaded trips last on average 2.5 weeks, with some variation within an origin-destination pair. Ballast trips last less, 1.74 weeks on average. Contracts are signed on average six days prior to the loading date.[9] Upon signing a contract, about 42% of ships are already in the loading port. Ships that do not find a cargo, remain in their current location with probability 77%. Clarksons reports the product carried in a small subsample of the contract dataset (about 20%). The main commodity categories are grain (23%), iron ore (20%), coal (20%), alumina ore (6%), chemicals/fertilizer (6%) and minor bulks like wood chips and sands. Finally, it is worth noting that our sample period is one of low shipping demand and severe ship oversupply due to high ship investment between 2005 and 2008 (see Kalouptsidi (2014, 2017)).

2.3 Facts

In this section, we present a number of novel facts: we document geographic patterns of trade through vessel movements, we explore the nature of trade costs (i.e. shipping prices) and how they correlate with trade imbalances and product values, and we discuss some descriptive evidence suggestive of search frictions. Our findings guide the model formulation in Section 1.4.

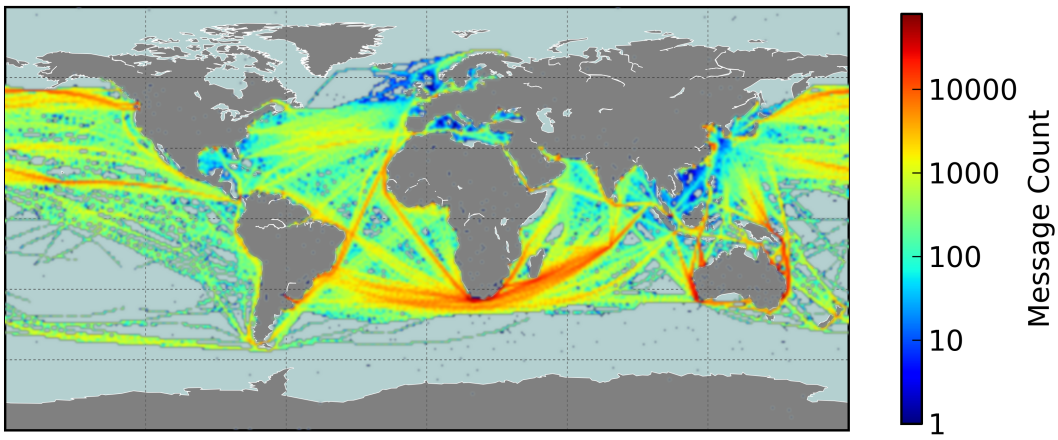

Figure 2.1: Number of AIS signals (messages) received in a ten day period.

2.3.1 Trade Flows, Geography and Ballasting

Figure 2.1 plots vessel locations during a 10 day period. It reveals that some of the most frequent voyages are between Australia and China, Brazil and China, as well as Northwest America and China. This graph does not distinguish between loaded and ballast voyages. The most popular loaded trip is from Australia to China (around 5% of loaded trips in our dataset). The most popular ballast trip is the reverse, from China to Australia (5.7% of ballast trips). It is no accident that China dominates the observed flows: in recent years, Chinese growth has led to massive imports of raw materials for industrial expansion and infrastructure building. In turn, Australia, Brazil and Northwest America are large exporters of minerals, grain, coal, etc.

China's example suggests that global trade in commodities features a substantial imbalance, partly owing to the natural inheritance of different countries in raw materials: we next illustrate that this feature is quite general and that most countries are either net importers or net exporters of commodities. Figure 2.2 plots the difference between the number of ships departing loaded and the number of ships arriving loaded,

[9]As practitioners say, "a ship is not a train"; it is not possible for a ship to promise too far in advance arrival to load at a specific port, due to the uncertainties of prior travels (weather, port/canal congestion, port strikes, etc.).

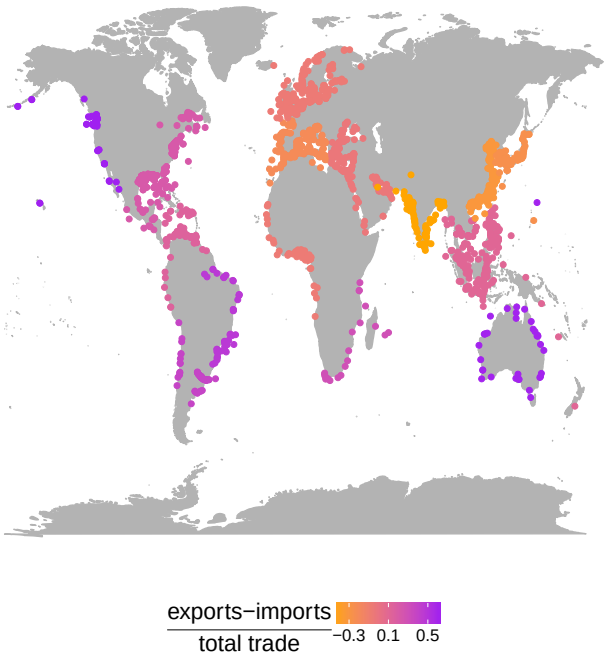

Figure 2.2: Trade Imbalance. Difference between exports (ships leaving loaded) and imports (ships arriving loaded) over total trade (all ships).

over the sum of the two. A positive ratio indicates that a country is a net exporter, while a negative ratio suggests that the country is a net importer; a ratio close to zero implies balanced trade. As shown, trade flows in most countries in the world are considerably imbalanced. Australia, Brazil and Northwest America are big exporters, whereas China and India are big importers. This feature of trade is not unique to raw materials; container shipping exhibits similar asymmetries, suggesting that trade imbalances may affect trade costs through a mechanism similar to the one discussed in the present paper; the direction of the imbalance, however, may be different (e.g. China is a big exporter in containers rather than a big importer).

A consequence of the imbalanced nature of international trade, is that ships spend much of their time traveling ballast, i.e. without cargo. We find that the fraction of the miles a ship travels ballast over the total miles traveled is about 45%.

The imbalanced nature of trade, although an important empirical feature, is often overlooked in the trade literature. In this paper, we do not assume balanced trade and in fact this asymmetry is a key driver of our model and empirical findings.

2.3.2 Trade Costs (Shipping Prices)

We next turn to the nature of trade costs that exporters face. A quick inspection of the data reveals that there are large asymmetries in trade costs across space: for instance, a trip from China to Australia costs on average 7,500 dollars per day, while a trip from Australia to China costs substantially more, at 10,000 dollars per day on average, net of fuel costs.[10] In fact, most trips exhibit substantial asymmetry: the average ratio of the price from origin i to destination j to the price from j to i (highest over lowest), is equal to 1.6 and can be as high as 4.1.[11] This empirical pattern suggests that bilateral distance is not the only determinant of trade costs.

To delve deeper into the determinants of trade costs, we present some price regressions in Table 2.2. The first column presents the results of a log-price regression on ship types and country of origin fixed effects which already account for 66% of the price variation, suggesting that geography is important in explaining trade costs. The second column adds destination fixed effects and, interestingly, the fraction of price variation explained increases. This suggests that ships may demand a premium to travel towards a destination with low exports, to compensate for the difficulty of finding a new cargo originating from that destination. To control more directly for this effect, in the third column of Table 2.2 we consider (i) the probability that ships leave ballast from the destination; and (ii) conditional on leaving ballast, the miles a ship travels ballast from the destination on average. As expected, we find that both variables are positive and significant and lead to substantially higher prices. Indeed, a 1% increase in the average distance traveled ballast after the destination, is associated with a 0.17% increase in prices. Similarly, exporting to a destination where

[10] This price asymmetry has been documented also in container shipping; see e.g. Wong (2017) and references therein.

[11] This is calculated using the 15 geographical regions employed in our empirical exercise below (see Section 1.6), to guarantee sufficient data per origin/destination.

the probability of a ballast trip afterwards is ten percentage points higher, costs on average 2.3% more.[12]

Finally, it is worth emphasizing that ship heterogeneity is not a first order issue. First, even in our small dataset, the majority of ships (80%) are seen carrying at least 2 of the 3 main products (coal, ore and grain), which suggests that ships do not specialize on certain products. Similarly, we observe that most ships travel to most regions, suggesting that they do not specialize geographically either. Ship fixed effects have no explanatory power in either price regressions or ballast probability regressions. Finally, this is consistent with much of the evidence provided in Kalouptsidi (2014); for instance hedonic regressions of ship resale prices suggest that unobserved heterogeneity is not an important consideration.

2.3.3 Search Frictions

In this section we present some descriptive evidence suggesting that search frictions inhibit the matching of all available ships and exporters. Overall, it is not straightforward to know a priori whether a market (here, the market for sea transport) suffers from search frictions. In labor markets, where search frictions are generally thought to be present, two main empirical regularities are often offered as evidence: (i) the coexistence of unemployed workers and vacant firms; and (ii) wage inequality among observationally identical workers. In this section we show that a similar set of empirical conditions hold in shipping, suggesting that search frictions may be present here as well.

We first consider whether there is evidence of unrealized matches, along the lines of (i). In particular, we cannot replicate the argument done for labor markets, since our data reports only ships and matches, not searching exporters (similar to vacancies

[12]To confirm that this result is not driven by the different composition of products exported towards different destinations, the last column of Table 2.2 also controls for the product carried for the subsample of contracts reporting this information. We discuss how trade costs depend on the value of the product in Section 2.3.3.

	I	II	III	IV
	log(price per day)			
Handymax	-0.148**	-0.136**	-0.123**	0.027
	(0.014)	(0.014)	(0.014)	(0.120)
Handysize	-0.397**	-0.330**	-0.343**	-0.209**
	(0.017)	(0.018)	(0.017)	(0.124)
Panamax	-0.223**	-0.214**	-0.212**	-0.117
	(0.013)	(0.013)	(0.013)	(0.119)
Coal				0.088**
				(0.045)
Fertilizer				0.245**
				(0.051)
Grain				0.131**
				(0.048)
Ore				0.124**
				(0.045)
Steel				0.135**
				(0.049)
Probability of ballast			0.234**	0.556**
			(0.030)	(0.081)
Average duration of ballast trip (log)			0.166**	0.065**
			(0.014)	(0.032)
Constant	10.304**	10.284**	9.127**	8.915**
	(0.068)	(0.103)	(0.099)	(0.408)
Destination FE	No	Yes	No	No
Origin FE	Yes	Yes	Yes	Yes
Product FE	No	No	No	Yes
Quarter FE	Yes	Yes	Yes	Yes
Obs	11,014	11,014	11,011	1,662
R^2	0.663	0.694	0.674	0.664

$^{**}p < 0.05$, $^{*}p < 0.1$

Table 2.2: Shipping price regressions.

in labor markets); we can, however, consider a different moment that has a similar flavor.

Figure 2.3 displays the weekly number of ships that arrive empty and load, as well as the number of ships that leave empty, in Norway and in Chile. Both countries are net exporters. In Norway, several ships arrive empty and load, but almost no ship departs empty. In Chile, however, the picture is quite different: we simultaneously see ships that arrive empty to load, and ships that depart empty. In other words, it frequently happens that an empty ship arrives and picks up cargo, while at the same time another ship departs empty. This is suggestive of wastefulness in Chile: why does the ship that depart empty, not pick up the cargo, instead of having another ship arrive from elsewhere to pick it up?

We next perform this exercise for all net exporting countries, by computing the bi-weekly ratio of the incoming empty and loading ships over the outgoing empty ships for a given country. In the absence of frictions, one may expect this ratio to be close to zero. However, as shown in Figure 2.4 which depicts the histogram of these ratios, most countries are more similar to Chile, than Norway. Indeed, the average ratio is well above zero and for some countries it is even above 0.5. In addition, this pattern is quite robust in a number of dimensions.[13] While in labor markets, as some researchers have argued, observed or unobserved heterogeneity may partly explain the co-existence of unemployment and vacancies (a vacancy for a chemical engineer may not be of interest to a high school dropout), in this market the importance of

[13]This figure is robust to alternative market definitions, time periods and ship types. Figure 2.14 in the Appendix presents this histogram by ship type: Capesize vessels exhibit somewhat larger mass towards zero, consistent with the somewhat higher concentration of ships and charterers, as well as the ships' ability to approach fewer ports. The figure is also the same if done by port rather than country. As mentioned above, ships tend to carry all products; thus we do not believe this pattern is explained by product switching. Labor contracts are usually about 5-8 months long and the crew flies between their home and the relevant port. To control for repairs we remove stops longer than 6 weeks. Finally, we only consider as "ships arriving empty" the ships arriving empty and sailing full toward another region, and we consider as "ships leaving empty" ships sailing empty toward a different country; so movements to nearby ports are excluded. This definition also implies that refueling cannot explain the fact either- though there are very small differences in fuel prices across space anyway (less than 10%).

heterogeneity is much more limited: as discussed above ships are widely considered to offer homogeneous services and do not specialize geographically or in terms of products.

Again inspired by the labor literature, we investigate a second aspect of this market that is suggestive of search frictions: dispersion in prices. In markets with no frictions, the law of one price holds, so that there is a single price for the same service. This does not hold in labor markets, where there is large wage dispersion among workers who are observationally identical. This observation has generated a substantial and widely influential literature on frictional wage inequality, i.e. wage inequality that is driven by search frictions.[14]

In the shipping market a similar empirical regularity is present. As we already saw in Table 2.2 there is substantial price dispersion in shipping contracts. More specifically, at best we can account for about 70% of price variation, controlling for ship size, as well as quarter, origin and destination fixed effects. Moreover, the coefficient of variation of prices within a given quarter, origin and destination triplet is about 30% (23%) on average (median).[15] For instance, in the most popular trip from Australia to China the coefficient of variation is on average 37% and ranges from 21% to 55% across quarters.

[14]See for instance Burdett and Mortensen (1998), Postel-Vinay and Robin (2002), Mortensen (2003) and references therein.

[15]Again here we use the 15 geographical regions employed in our empirical exercise below (see Section 1.6), to guarantee sufficient data per origin/destination.

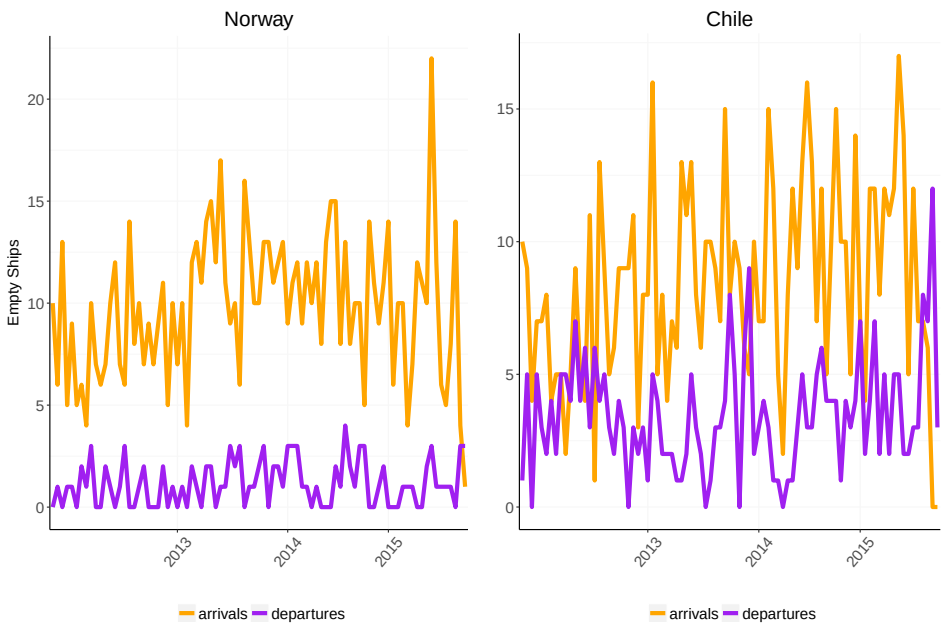

Figure 2.3: Flow of ships arriving empty and loading, and ships leaving empty in 2 week intervals.

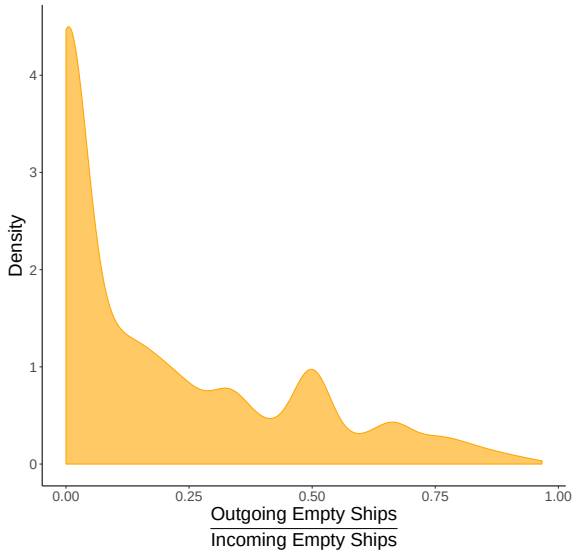

Figure 2.4: Histogram of the ratio of outgoing empty over incoming empty and loading ships in net exporting countries.

In addition, it is worth noting that the type of product carried affects the price paid and overall more valuable goods lead to higher contracted prices. In the absence of frictions, if there are more ships than exporters, as is arguably the case during our

sample period, we would expect prices to be bid down to the ships' opportunity cost.[16] In contrast, in markets with frictions and bilateral bargaining, as shown formally in the model of Section 1.4, the buyer's valuation affects the price he pays. In the context of the shipping market, our model predicts that exporters with higher valuations pay more when there are search frictions, consistent with the evidence in Table 2.2.

Finally, we perform a simple "dispatcher" simulation, which assigns every observed match to the ship that is closest to it. We find that the fraction of distance a ship ballasts in the simulation is 38% on average, significantly less than in the data (45%). In this simulation, ships do not optimally choose where to search; we return to this in the counterfactuals of Section 2.7.

We revisit search frictions in Section 2.5.1, where we both provide a test, as well as estimate a non-parametric matching function, thus flexibly measuring the extent of search frictions. We close this section by briefly discussing what the nature of search frictions may be. The mere existence of brokers suggests that search frictions are an issue in this market. Information frictions may still prevail though; when a ship is unmatched in a certain geographical region (e.g. East Americas) her broker may not "meet" the broker of a specific exporter in one of the ports. Interviews with industry participants indeed suggest that information can get lost between the different brokers of different shipowners and exporters; interestingly, oftentimes more than two brokers may mediate the deal. Unfortunately direct data on brokerage arrangements are not available. In a market with a small number of large exporters, it might be easier to for them to be matched to the existing ships. Consistent with this, we find that price dispersion is negatively correlated with the Herfindahl Index of the observed ship charterers.

[16] In a frictionless market with more ships than freights and homogeneous ships, in equilibrium the price from an origin to a destination would be such that ships are indifferent between transporting the cargo and staying unmatched.

2.4 Model

We next introduce a dynamic spatial search model of the global shipping industry. Geography enters the model through different trip durations across different locations. There are two types of agents: exporters and ships. Exporters choose whether and where to export, while unmatched ships choose where to ballast. Following the search and matching literature we model new matches every period using a matching function, which captures the implications of frictional trading, in a parsimonious fashion (Pissarides (2000)). In other words, we do not explicitly model the meeting technology between exporters and ships, but given the evidence presented in Section 2.3.3 we allow for the possibility of search frictions.

2.4.1 Environment

Time is discrete. There are I locations/regions, $i \in \{1, 2, ..., I\}$. There two types of agents, exporters and ships. Both are risk neutral and have discount factor β.

Exporters At each location i and period t, there are f_{it} freights that need to be delivered to another location. We use the words freight, exporter and cargo interchangeably. Exporters have heterogeneous valuations, v, from exporting to their destination. The valuation of a freight going from i to j is drawn from the distribution F_{ij}^{v} with mean μ_{ij}.[17] Unmatched freights survive with probability $\delta > 0$. Every period, at each location i, $\mathcal{E}_i$ potential exporters decide whether and where to export. If they decide to export, they pay production and export costs, κ_{ij} and draw their valuation v.

[17] In this paper, we do not consider the determination of commodity prices; in other words, we take exporter valuations to be exogenous. v is meant to capture the exporter's revenue. Determining this object in equilibrium requires additional data on exporters and is an interesting avenue for future research.

Ships There are S homogeneous ships in the world.[18,19] In every period, a ship is either traveling loaded or ballast, from some location i to some location j, or it is at port in some region i. A ship at port in location i incurs a per period cost c_i^u, while a ship sailing from i to j incurs a per period cost c_{ij}^s. The duration of a trip between region i and region j is stochastic: a traveling ship arrives at j this period with probability ξ_{ij}, so that the average duration of the trip is $1/\xi_{ij}$.[20]

Matching Freights can only be delivered to their destination by ships and each ship can carry (at most) one freight. As discussed earlier, we capture the process through which exporters are paired with available ships in a market by a matching function, whereby the number of matches at time t are

$$m_{it} = m_i\left(s_{it}, f_{it}\right)$$

where s_{it} is the number of unmatched ships in i. The matching function $m_i\left(s_{it}, f_{it}\right)$ is increasing in both arguments, and whenever $m_{it} = \min\left(s_{it}, f_{it}\right)$ the matching process is frictionless. Let λ_{it} denote the probability with which an unmatched ship in location i meets a freight; $\lambda_{it} = m_{it}/s_{it}$. Similarly, let λ_{it}^f denote the probability with which an unmatched freight meets a ship; $\lambda_{it}^f = m_{it}/f_{it}$.

When a ship and a freight meet, they can either agree on a price to be paid by the freight to the ship or they both revert to their outside options. Note that the outside option of a freight is to remain unmatched and wait for another ship, while

[18] We follow Kalouptsidi (2014) and assume constant returns to scale so that a shipowner is a ship. Similarly, a freight owner is a freight, so that he does not choose the export tonnage. We also ignore the different ship sizes in the model; estimation results are robust when we consider types separately.

[19] In this paper, we do not model ship entry and exit. Exit is overall very small, while due to long construction lags in shipbuilding (two to six years), the fleet is fixed in the short run; see Kalouptsidi (2014, 2017).

[20] It is straightforward to have deterministic trip durations instead. Our specification, however, preserves tractability and allows for some variability e.g. due to weather shocks, without affecting the steady state properties of the model.

the outside option of the ship is to either remain unmatched in the current region or to ballast elsewhere. The surplus of the match over the parties' outside options is split between the freight and the ship via the price-setting mechanism. We assume that the price, τ_{ijv}, paid to the ship delivering a freight of valuation v from region i to destination j, is determined by generalized Nash bargaining, with $\gamma \in (0,1)$ denoting the exporter's bargaining power. The price is paid upfront and the ship commits to begin its voyage immediately to region j.

In what follows, we assume that exporter valuations are sufficiently high so that in equilibrium, when a ship and an exporter meet, they always agree to form a match.[21] The agreed upon price guarantees that the exporter prefers not to wait for the following period (recall that ships are homogeneous) and that the ship does not gain from either waiting or ballasting elsewhere.

Finally, it is worth noting that matching occurs only if ships and freights are in the same region. This assumption is consistent with our data, where most ships are already in the region of the freight when signing a contract; contracts are signed only six days before loading; and there is substantial ship excess supply in all regions. In our estimation, each "region" is fairly broad usually comprising of several countries.

Timing The timing of each period is as follows:

1. In each region i, existing ships and exporters match.

2. In each region i, ships at port pay the port costs, draw additive iid preference shocks $\epsilon = [\epsilon_1, ..., \epsilon_I] \in \mathbb{R}^I$ from distribution F^ϵ and decide whether to (i) stay in

[21] An exporter (v, j) either forms a match with any ship it meets as long as his value, v, is high enough to generate a positive surplus given the destination, j, or cannot agree to a mutually acceptable price with any ship. Given that in this context valuations are an order of magnitude greater than the shipping price, this assumption is fairly innocuous (see Section 2.5.3).

their current region and wait for freight; or (ii) ballast toward some destination j.[22]

3. Unmatched ships that decided to ballast away begin traveling to their chosen destination. All ships already traveling from i to j arrive at j with probability ξ_{ij}. Existing exporters disappear with probability $1-\delta$.

4. In each region i, $\mathcal{E}_i$ potential exporters decide whether and to which destination to export to. The exporters that do enter the market draw their valuations from F_{ij}^v, pay cost κ_{ij} and join the pool of unmatched exporters the following period.

States and Transitions The state variable of a ship in region i includes its current location i, as well as the vector (s_t, f_t, s_t^w) where $s_t = [s_{1t}, ..., s_{It}]$, $f_t = [f_{1t}, ..., f_{It}]$ and the $I^2 - I$ dimensional vector s_t^w, with entries s_{ijt}^w, denotes the number of ships traveling from i to j in period t. The state variable of an existing exporter in i includes his location i, valuation v and destination j, as well as the vector (s_t, f_t, s_t^w). Exporters in region i at time t transition as follows:

$$f_{it+1} = \delta \left(f_{it} - m_i \left(s_{it}, f_{it} \right) \right) + d_i \tag{2.1}$$

with d_i the (endogenous) flow of new freights. Ships at location i transition as follows:

$$s_{it+1} = \left(s_{it} - m_i \left(s_{it}, f_{it} \right) \right) P_{ii} + \sum_{j \neq i} \xi_{ji} s_{jit}^w \tag{2.2}$$

where P_{ij} is the probability of an unmatched ship ballasting from i to j (determined endogenously, see below). In words, out of s_{it} ships, m_{it} ships get matched and leave

[22] Recall that the average trip duration is 2-3 weeks, while contracts are signed on average six days in advance, so that ships often sail towards a destination without having signed a contract already. This is fairly intuitive given the large trade imbalances and high port costs.

i, while out of the ships that did not find a match, fraction P_{ii} chooses to remain at i rather than ballast away; moreover, out of the ships traveling towards i, fraction ξ_{ji} arrive. Finally, ships that are traveling from i to j, s^w_{ijt} evolve as follows:

$$s^w_{ijt+1} = (1 - \xi_{ij})\, s^w_{ijt} + P_{ij}\,(s_{it} - m_i\,(s_{it}, f_{it})) + G_{ij} m_i\,(s_{it}, f_{it}) \tag{2.3}$$

where G_{ij} is the probability of going loaded from i to j (determined endogenously, see below). In words, fraction ξ_{ij} of the traveling ships arrive, fraction P_{ij} of ships that remained unmatched in location i chose to ballast to j and finally, G_{ij} of ships matched in i depart loaded to j.

2.4.2 Equilibrium

We derive the optimal behavior of exporters and ships, as well as the equilibrium prices and trade flows. In this paper, we consider the steady state of this model.[23] This assumption is not unreasonable for the data at hand, which covers a period (2012-2016) that is uniformly characterized by ship oversupply and relatively low demand for shipping services. More specifically, we assume that agents view the spatial distribution of ships and freights, (s_t, f_t, s^w_t), as fixed and make decisions based on their steady-state values: given the short-lived nature of their decisions (where to ballast) it does not feel unreasonable that they ignore aggregate long-run shocks when making these weekly choices.

Ships Let W_{ij} denote the value of a ship traveling from i to j (empty or loaded). Then:

$$W_{ij} = -c^s_{ij} + \xi_{ij}\beta U_j + (1 - \xi_{ij})\beta W_{ij} \tag{2.4}$$

[23] Given that a ship can travel to most ports in the world in under a month, any transition dynamics to a new steady state will be short. This is convenient in our counterfactual analysis of Section 2.7, where we are able to compare steady states directly.

In words, the ship that is traveling from i to j, pays per period cost of transit c^s_{ij}; with probability ξ_{ij} it arrives at its destination j, where it will begin unmatched with value U_j defined below; finally, with the remaining probability $1 - \xi_{ij}$ the ship does not arrive and keeps traveling. The ship arrives at j after $1/\xi_{ij}$ periods on average.

Consider now a ship in region i. This ship obtains:

$$U_i = -c^u_i + \lambda_i E_{j,v} V_{ijv} + (1 - \lambda_i) J_i \tag{2.5}$$

In words, the ship is matched with probability λ_i, in which case it obtains the value of a matched ship V_{ijv} defined below. The ship takes expectation over the type of freight it meets, i.e. its value and destination. With the remaining probability, $1-\lambda_i$, the ship does not find a freight and it obtains the value J_i also defined below. Finally, the ship pays the per period port cost c^u_i.

If matched with an exporter with value v and destination j, the ship receives the agreed upon price, τ_{ijv}, and begins traveling, so that:

$$V_{ijv} = \tau_{ijv} + W_{ij} \tag{2.6}$$

If the ship remains unmatched, it faces the choice of either staying at i and matching there the following period with probability λ_i, or ballasting away from i in search of better opportunities. In the latter case, the ship can choose among all possible destinations. In particular, if unmatched, the ship draws preference shocks $\epsilon \in \mathbb{R}^I$, from a double exponential distribution F^ϵ, with standard deviation σ. The unmatched ship's value function is:

$$J_i(\epsilon) = \max \left\{ \beta U_i + \sigma \epsilon_i, \max_{j \neq i} W_{ij} + \sigma \epsilon_j \right\} \tag{2.7}$$

and let:

$$J_i \equiv E_\epsilon J_i(\epsilon) = \sigma \log \left(\exp \frac{\beta U_i}{\sigma} + \sum_{j \neq i} \exp \frac{W_{ij}}{\sigma} \right) + \sigma \gamma^{euler}$$

where γ^{euler} is the Euler constant.[24] In words, if the ship stays in its current region i, it obtains value U_i; otherwise the ship chooses another region and begins its trip there. Let P_{ii} denote the probability that a ship in location i chooses to remain there, while P_{ij} denote the probability that a ship at i chooses to ballast to j. We have:

$$P_{ii} = \frac{\exp\left(\beta U_i/\sigma\right)}{\exp\left(\beta U_i/\sigma\right) + \sum_{l \neq i} \exp\left(W_{il}/\sigma\right)} \tag{2.8}$$

and

$$P_{ij} = \frac{\exp\left(W_{ij}/\sigma\right)}{\exp\left(\beta U_i/\sigma\right) + \sum_{l \neq i} \exp\left(W_{il}/\sigma\right)}. \tag{2.9}$$

Exporters We start with existing exporters and then consider exporter entry. An exporter that is matched in location i receives value:

$$V^f_{ijv} = v - \tau_{ijv}, \tag{2.10}$$

in words, he obtains his delivery value, v and pays the agreed price, τ_{ijv}, which is derived below. An exporter that does not get matched receives no payoff in the period and survives with probability δ; if so, the following period with probability λ^f_i he gets matched and receives V^f_{ijv}, while with the remaining probability $1 - \lambda^f_i$ he remains

[24]The formula for the ex ante value function $J_i = E_\epsilon J_i(\epsilon)$ is the closed form expression for the expectation of the maximum over multiple choices, and is obtained by integrating $J_i(\epsilon)$ over the double exponential distribution of ϵ.

unmatched again:

$$J_{ijv}^{f} = \beta\delta\left(\lambda_i^f V_{ijv}^{f} + \left(1-\lambda_i^f\right) J_{ijv}^{f}\right). \tag{2.11}$$

There are $\mathcal{E}_i$ ex ante homogeneous potential exporters in location i in every period. Each potential entrant in region i, makes a discrete choice between not exporting, as well as which destination j to export to, subject to production and exporting costs κ_{ij}, as well as random preference shocks $\epsilon^f \in \mathbb{R}^I$, distributed according to a double exponential distribution. Upon deciding to become an existing exporter in i with destination j, the entrant draws valuation v from F_{ij}^{v}. Therefore, a potential entrant solves:

$$J_i^{ef} = \max\left\{\epsilon_0^f, \quad \max_{j\neq i}\left\{E_v J_{ijv}^{f} - \kappa_{ij} + \epsilon_j^f\right\}\right\}$$

where we denote by 0 the (outside) option of not exporting and normalize the payoff in that case to zero.[25]

Potential exporters' behavior is given by the choice probabilities:

$$\tilde{G}_{ij} \equiv \frac{\exp\left(J_{ij}^{f} - \kappa_{ij}\right)}{1+\sum_{l\neq i}\exp\left(J_{il}^{f} - \kappa_{il}\right)} \tag{2.12}$$

and

$$\tilde{G}_{i0} \equiv \frac{1}{1+\sum_{l\neq i}\exp\left(J_{il}^{f} - \kappa_{il}\right)} \tag{2.13}$$

where $J_{ij}^{f} \equiv E_v J_{ijv}^{f}$.

[25]It is possible to allow the potential exporters to know their valuations (across destinations) before making their exporting choice, but this would make the estimation computationally more demanding.

Therefore the number of entrant exporters in i equals:

$$d_i = \mathcal{E}_i \left(1 - \tilde{G}_{i0}\right) \tag{2.14}$$

It is worth noting that the distribution of export destinations conditional on exporting is given by

$$G_{ij} \equiv \frac{\tilde{G}_{ij}}{1 - \tilde{G}_{i0}} \tag{2.15}$$

This is the distribution that ships employ when forming expectations over the potential matches in different regions in equation (2.5).

Prices As discussed above, the rents generated by a match between a freight and a ship, are split via Nash bargaining. This implies the usual surplus sharing condition:

$$\gamma \left(V_{ijv} - J_i\right) = (1 - \gamma) \left(V^f_{ijv} - J^f_{ijv}\right) \tag{2.16}$$

We use the above condition to solve out for the equilibrium price τ_{ijv}, in the following lemma: The agreed upon price between a ship and an exporter with valuation v and destination j in location i is given by:

$$\tau_{ijv} = \frac{\gamma \left(1 - \beta\delta \left(1 - \lambda^f_i\right)\right)}{1 - \beta\delta \left(1 - \gamma\lambda^f_i\right)} \left(J_i - W_{ij}\right) + \frac{(1 - \gamma)\left(1 - \beta\delta\right)}{1 - \beta\delta \left(1 - \gamma\lambda^f_i\right)} v \tag{2.17}$$

Substitute V_{ijv}, V^f_{ijv}, J^f_{ijv} and U_i in (2.16). In other words, the price is a linear combination of the exporter's value, v, and the difference between the ship's value of transporting the freight, W_{ij}, and its outside option, J_i.

It is worth noting that the price is decreasing in the value of a ship traveling from i to j, W_{ij}. Recall that W_{ij} depends on both the conditions at the destination through U_j, and distance, captured by ξ_{ij}, since $W_{ij} = -c^s_{ij}/(1-\beta(1-\xi_{ij})) + \xi_{ij}\beta U_j/(1-\beta(1-\xi_{ij}))$. In other words, destinations that are unappealing to ships because there are few freights there and the probability of ballasting afterwards is high, would command higher prices. This is consistent with the evidence presented in Table 2.2. The same holds for destinations that are further away (low ξ_{ij}), have low value freights or severe search frictions. Moreover, U_j also controls for conditions at all possible ballast destinations from j, as well as for conditions at all possible export destinations from j. Finally, note that the price between i and j depends on all countries rather than just i and j, both through the outside option of the ship, J_i, as well as through the conditions in the export/ballast destinations from j, captured in U_j, as discussed above. We return to the importance of the network when we perform model counterfactuals in Section 2.7.

In addition, exporters that have a higher value, v, pay higher prices, again consistent with evidence in Table 2.2. As discussed in Section 2.3.3, this is true because the law of one price no longer holds when there are search frictions. In a world without search frictions and more ships than freights, the shipping price is given by $\tau_{ij} = J_i - W_{ij}$; therefore the previous properties of trade costs (dependence on distance, destination and entire network of countries) are independent of the presence of search frictions and *still hold*.[26]

Steady State Equilibrium We next define the steady state equilibrium for this model and prove that it exists. A steady state equilibrium, (s^*, f^*, s^{w*}), is a distribution of ships and exporters over locations, that satisfies the following conditions:
(i) Ships' optimal behavior, $P_{ij}(s^*, f^*, s^{w*})$ follows (2.8) and (2.9) and expectations

[26]Since ships are on the "long side" of the market, the price has to be such that ships are indifferent between loading and going to destination j and remaining unmatched, i.e. it must be that $V_{ijv} = J_i$. Substituting in for V_{ijv} from equation (2.6) leads to the above price equation.

employ (2.15).

(ii) Potential exporters' behavior, $\tilde{G}_{ij}\left(s^*, f^*, s^{w*}\right)$, follows (2.12) and (2.13) and entrants are determined from (2.14).

(iii) Prices are determined by Nash bargaining, according to (2.17).

(iv) Ships and freights satisfy the following steady state equations (established in Proposition 2.4.2 below):

$$s_i^* = \sum_j P_{ji}\left(s^*, f^*, s^{w*}\right)\left(s_j^* - m_j\left(s_j^*, f_j^*\right)\right) + \sum_{j \neq i} G_{ji}\left(s^*, f^*, s^{w*}\right) m_j\left(s_j^*, f_j^*\right) \quad (2.18)$$

$$f_i^* = \delta\left(f_i^* - m_i\left(s_i^*, f_i^*\right)\right) + E_i\left(1 - \tilde{G}_{i0}\left(s^*, f^*, s^{w*}\right)\right) \quad (2.19)$$

$$s_{ij}^{w*} = \frac{1}{\xi_{ij}}\left(P_{ij}\left(s^*, f^*, s^{w*}\right)\left(s_i^* - m_i\left(s_i^*, f_i^*\right)\right) + G_{ij}\left(s^*, f^*, s^{w*}\right) m_i\left(s_i^*, f_i^*\right)\right)$$

Suppose that the matching function is continuous, the preference shocks ϵ and ϵ^f have full support, $\mathcal{E}_i$ and S are finite and $f_i \leq \mathcal{E}_i/(1-\delta)$. Then, an equilibrium exists. See the Appendix.

Trade Flows Finally, we characterize the steady state trade flow between regions i and j, which is equal to:

$$\mathcal{E}_i \tilde{G}_{ij} = \mathcal{E}_i \frac{\exp\left(J_{ij}^f - \kappa_{ij}\right)}{1 + \sum_{l \neq i} \exp\left(J_{il}^f - \kappa_{il}\right)} = \mathcal{E}_i \frac{\exp\left(\alpha_i\left(\mu_{ij} - \tau_{ij}\right) - \kappa_{ij}\right)}{1 + \sum_{l \neq i} \exp\left(\alpha_i\left(\mu_{il} - \tau_{il}\right) - \kappa_{il}\right)}$$

where $\alpha_i = \beta\delta\lambda_i^f / \left(1 - \beta\delta\left(1 - \lambda_i^f\right)\right)$, since in the steady state:

$$J_{ij}^f \equiv E_v J_{ijv}^f = E_v \frac{\beta\delta\lambda_i^f\left(v - \tau_{ijv}\right)}{1 - \beta\delta\left(1 - \lambda_i^f\right)} = \alpha_i\left(\mu_{ij} - \tau_{ij}\right) \quad (2.20)$$

where $\tau_{ij} \equiv E_v \tau_{ijv}$. This equation is reminiscent of a "gravity equation"; it delivers the trade flow (in quantity rather than value) from i to j as a function of two components. First, the primitives $\{\lambda_i^f, \mu_{ij}, \kappa_{ij}, \mathcal{E}_i\}$ not just for i and j but for all regions; this is reminiscent of the analysis in Anderson and Van Wincoop (2003) who show that the gravity equation in a trade model needs to include a country's overall trade disposition. Second, the endogenous trade costs, τ_{ij}, for all j. The important addition here is that the trade flow depends on all countries through the endogenous trade cost τ_{ij}; indeed, recall that τ_{ij} depends on all locations both through the outside option of the ship that can ballast anywhere, but also because the ship cares about the conditions in the ballast and export destinations from location j. Therefore, any change in the primitives affects trade flows both directly, but also indirectly through its impact on trade costs. We illustrate this when we perform model counterfactuals in Section 2.7.

2.5 Empirical Strategy

In this section we lay out the empirical strategy followed to estimate the model of Section 2.4. The main model primitives we wish to recover are: the matching function and searching exporters, the ship travel and port costs, as well as the distribution of exporter valuations and their costs. We describe the empirical strategy here, while in Section 2.6 we present the results.

2.5.1 Matching Function Estimation

A sizable literature has estimated matching functions in several different contexts (e.g. labor markets, marriage markets, taxicabs). For instance, in labor markets, one can use data on unemployed workers, job vacancies and matches to recover the underlying matching function. In the market for taxi rides, one observes taxis and their rides,

but not hailing passengers; in recent work, Buchholz (2016), and Frechette et al. (2016) have used such data, coupled with a "parametric" assumption on the matching function to recover the hailing passengers.[27]

Similar to the taxi market, we observe ships and matches, but not searching exporters. Here, we adopt a novel approach to simultaneously recover both exporters, as well as a nonparametric matching function. Our approach makes two contributions to the literature. First, we do not take a stance on the presence and magnitude of search frictions in the industry. Why is this important? Consider the case of no search frictions, so that the matching function is

$$m_{it} = \min\left(s_{it}, f_{it}\right) \tag{2.21}$$

In words, all potential matches are realized. In contrast, if there are search frictions, we have:

$$m_{it} = m_i\left(s_{it}, f_{it}\right) \leq \min\left(s_{it}, f_{it}\right) \tag{2.22}$$

so that some potential matches are not realized. If one side of the market is unobserved (here f_{it}) or mismeasured (arguably, in labor markets searching workers are imperfectly measured; Lange and Papageorgiou (2017)) it is not straightforward to differentiate (2.21) from (2.22). Indeed, when a ship/taxi is traveling empty is it because no exporter/passenger was searching or because an exporter/passenger was there but did not get to meet the ship/taxi due to frictions? Our approach, allows us to disentangle the two.

[27]Buchholz (2016)assumes an "urn-ball" matching function. Frechette et al. (2016) construct a numerical simulation of taxi drivers that randomly meet passengers over a grid that resembles Manhattan; this spatial simulation essentially corresponds to the matching function, and can be inverted to recover hailing passengers.

Our second contribution is to avoid imposing parametric restrictions on the matching function. The literature has imposed functional forms such as the Cobb-Douglas. The desire to be non-parametric is not just "stylistic" when it comes to matching functions: parametric restrictions are directly linked to welfare implications. For instance, it has been shown in a wide class of labor market models, that the condition for constrained efficiency depends crucially on the elasticity of the matching function with respect to the search input (Hosios (1990)). In much of the matching function estimation literature this elasticity has been restricted to be constant.

Our approach borrows from the literature on nonparametric identification (Matzkin (2003)). Roughly, the method leverages (i) an invertibility assumption between matches and freights, (ii) the observed relationship between ships and matches, (iii) an instrument that shifts the number of ships, and (iv) a restriction on the matching function that allows us to disentangle monotonic transformations. To provide some intuition, we outline a simple version of the methodology. We then formalize the argument. We refer the interested reader to Matzkin (2003) for further details.

Suppose that (i) the matching function $m_i(s_i, f_i)$ is continuous and strictly increasing in f_i and, (ii) that s_i is independent of f_i. The first assumption is natural in our context: more freights should lead to more matches, all else equal. The second assumption will prove useful in presenting the estimation methodology, but is likely not valid in our case, as the spatial distribution of ships and freights is determined jointly in equilibrium; we relax this assumption below.

Let $F_{m|s}$ denote the distribution of matches conditional on ships, and F_f the distribution of freights, f. Then at a given point (s_{it}, f_{it}, m_{it}) we have:

$$
\begin{aligned}
F_{m|s}\left(m_{it}|s_{it}\right) &= \Pr\left(m_i\left(s,f\right) \leq m_{it}|s_{it}\right) \\
\text{monotonicity} &= \Pr\left(f \leq m_i^{-1}\left(s,m_{it}\right)|s_{it}\right) \\
\text{independence} &= \Pr\left(f \leq m_i^{-1}\left(s_{it},m_{it}\right)\right) \\
&= F_f\left(f_{it}\right) \qquad (2.23)
\end{aligned}
$$

In words, the conditional distribution of matches (outcome) on ships (observed covariate) at a point (m_{it}, s_{it}) is equal to the distribution of freights at the corresponding (unobserved) point f_{it}. Equation (2.23) is our main relationship for the identification and estimation of both freights and the matching function. However, (2.23) alone is not sufficient: it is not possible to distinguish monotonic transformations of f and $m(\cdot)$. To do so, a restriction on either the distribution F_f or the matching function is required. In this paper we assume that the matching function is homogeneous of degree one, so that: $m_i(\alpha s_i, \alpha f_i) = \alpha m_i(s_i, f_i)$, all $\alpha > 0$.[28] The intuition behind the identification argument is as follows: the correlation between s_i and m_i informs us on $\partial m_i(s_i, f_i)/\partial s_i$, since the sensitivity of matches to changes in ships is observed and s_i is independent of f_i by assumption; then, due to homogeneity, this derivative also delivers the derivative $\partial m_i(s_i, f_i)/\partial f_i$; and once these derivatives are known, so is the matching function, which can now be inverted to provide the freights as well.[29]

Finally, as mentioned above, independence of ships and freights is not a natural assumption in our setting. We do, however, have a plausibly valid instrument: sea

[28] In the labor literature, most matching function estimates find support for constant returns to scale (see Petrongolo and Pissarides (2001)). Given that the nature of search frictions is not necessarily that different (in both cases it is a shortcut for information frictions about which ships/freights may be available), we consider this a natural starting point.

[29] We could alternatively impose an assumption on the distribution F_f. For example, if we assume that F_f is uniform on $[0, 1]$, we can use (2.23) to recover f_{it} pointwise from the conditional distribution of m on s; once freights are recovered, we also instantly know the (inverse) matching function. Bajari and Benkard (2005) employ this methodology to nonparametrically estimate hedonic price equations and unobserved product quality in the case of personal computers.

weather shifts the arrival of ships in a port without affecting the number of freights. We therefore assume that ships s are a function of the instrument z and a shock η, such that ships and freights are independent conditional on η. This allows us to modify (2.23) by conditioning on $\hat{\eta}$ as well as on s to obtain the distribution of freights.

Proposition 2.5.1 formalizes these arguments: (i) Suppose that $m(s, f)$ is continuous, (positively) homogeneous of degree 1 and strictly increasing in f. Suppose further that s and f are independent. Finally, suppose that m is known for a specific pair (s^*, f^*) so that $m^* = m(s^*, f^*)$, with $m^* \neq 0$. Then, the function $m(\cdot)$ is identified. (ii) Suppose there exists an instrument z such that

$$s = h(z, \eta) \tag{2.24}$$

with z independent from (f, η). Assume that proper conditions hold so that η can be uniquely recovered from (2.24). Then, s and f are conditionally independent given η. The distribution F_f and the matching function can be recovered from:

$$F_{f|\eta}(\phi) = F_{m|s=\frac{\phi}{f^*}s^*,\eta}\left(\frac{\phi}{f^*}m^*\right)$$

$$F_f(\phi) = \int F_{f|\eta}(\phi) f_\eta(\eta) d\eta$$

$$m(s, \phi) = \int F^{-1}_{m|s,\eta}(F_{f|\eta}(\phi)) f_\eta(\eta) d\eta$$

The matching function is estimated separately for each region i; Section 2.6 presents the results.[30,31]

2.5.2 Ship Costs

We next turn to the ships' sailing costs, c_{ij}^s, and port costs, c_i^u, as well as the standard deviation of the shocks, σ, which is identified because the observed prices pin down the scale of payoffs (in dollars). We obtain the parameters of interest, $\theta = \{c_{ij}^s, c_j^u, \sigma\}$, from the ships' optimal ballast choice probabilities, (2.8) and (2.9), which are a function of the ships' value functions (U_i,W_{ij}), which in turn depend on the parameters of interest, θ. We estimate θ via Maximum Likelihood. We use a nested fixed point algorithm to solve for the ship value functions at every guess of the parameter values (Rust (1987)), compute the predicted choice probabilities and then calculate the likelihood. We calibrate the discount factor to $\beta = 0.995$.

Since our model features a number of inter-related value functions (W, U, V), it does not fall strictly into the standard Bellman formulation. Hence, we provide Lemma 2.5.2, which proves that our problem is characterized by a contraction map and thus the value functions are well defined. For each value of the parameter vector θ, the map $T_\theta : \mathbb{R}^I \to \mathbb{R}^I$, $U \to T_\theta(U)$ with,

$$T_\theta(U)$$
$$= -c_i^u + \lambda_i \sum_{j \neq i} G_{ij}\tau_{ij} + \lambda_i \sum_{j \neq i} G_{ij} \left[-\frac{c_{ij}^s}{1 - \beta\left(1 - \xi_{ij}\right)} + \beta\xi_{ij} \frac{U_j}{1 - \beta\left(1 - \xi_{ij}\right)} \right] + \left(1 - \lambda_i\right) J_i(\theta, U)$$

[30]Following Proposition 2.5.1, we need a known point (m^*, s^*, f^*), such that $m(\alpha s^*, \alpha f^*) = \alpha m^*$. We choose $1 = m(s^*, 1)$, so that one exporter is always matched when there are s^* ships. We set s^* such that in all locations $m_i \leq f_i$ and we iterate over s^*. Note that this approach delivers a conservative bound on search frictions, since in principle we could allow for higher levels of exporters.

[31]We interpret the observed time-series variation as driven by short-run deviations from the steady state values.

where $\tau_{ij} = E_v \tau_{ijv}$ is the mean price from i to j, is a contraction and $U(\theta)$ is the unique fixed point. See the Appendix. In brief, our estimation algorithm proceeds in the following steps:

1. Guess an initial set of parameters $\{c_{ij}^s, c_i^u, \sigma\}$.

2. Solve for the ship value functions via a fixed point. Set an initial value U^0. Then at each iteration m and until convergence:

   (a) Solve for W^m from:

   $$W_{ij}^m = \frac{-c_{ij}^s + \xi_{ij}\beta U_j^m}{1 - \beta\left(1 - \xi_{ij}\right)}$$

   (b) Update J^m from:

   $$J_i^m = \sigma \log\left(\exp\frac{\beta U_i^m}{\sigma} + \sum_{j \neq i} \exp\frac{W_{ij}^m}{\sigma}\right) + \sigma\gamma^{euler}$$

   (c) Update U^{m+1} from:

   $$U_i^{m+1} = -c_i^u + \lambda_i E_{j,v}\tau_{ijv} + \lambda_i \sum_{j \neq i} G_{ij} W_{ij}^m + \left(1 - \lambda_i\right) J_j^m$$

   where we use the actual average prices from i to j, i.e., $E_{j,v}\tau_{ijv} = \sum_{j \neq i} G_{ij}\tau_{ij}$. Note that λ_i is known (it is simply the average ratio $\frac{1}{T}\sum m_{it}/s_{it}$). Similarly, G_{ij}, the probability that an exporter ships from i to j (conditional on exporting), is obtained directly from the observed trade flows (see Section 2.5.3).

3. Form the likelihood using the choice probabilities:

$$\mathcal{L} = \sum_i \sum_j \sum_l \sum_t y_{ijlt} \log P_{ij}(\theta) = \sum_i \sum_j \log P_{ij}(\theta)^{n_{ij}} \qquad (2.25)$$

where y_{ijm} is an indicator equal to 1 if ship l chose to go from i to j in week t, n_{ij} is the number of observations (ship-weeks) that we observe a ship in i choosing j, and $P_{ij}(\theta)$ are given by (2.8) and (2.9).[32]

Identification As is always the case in dynamic discrete choice models, not all parameters are identified and some restriction needs to be imposed. Here, we have $I^2 + 1$ parameters and $I^2 - I$ choice probabilities, so we require $I + 1$ restrictions; we show this formally, borrowing from the analysis of Kalouptsidi et al. (2016) in the Appendix. The additional restrictions amount to using the observed fuel price to determine c^s_{ij} for some i, j; see Section 2.6.2. We estimate the port costs $c^u_1, ..., c^u_I$, which may be capturing heterogeneous costs difficult to measure (actual port costs, ability to wait outside of port, etc.), as well as σ. We present our results in Section 2.6.

2.5.3 Exporter Valuations and Costs

Using the observed shipping prices, we back out exporter valuations in a straightforward manner. Indeed, consider the equilibrium price (2.17) solved with respect to the exporter's valuation:

$$v = \frac{1 - \beta\delta\left(1 - \gamma\lambda^f_i\right)}{(1-\gamma)(1-\beta\delta)}\tau_{ijv} - \frac{\gamma\left(1 - \beta\delta\left(1 - \lambda^f_i\right)\right)}{(1-\gamma)(1-\beta\delta)}(J_i - W_{ij}) \tag{2.26}$$

Note that once $\theta = \{c^s_{ij}, c^u_i, \sigma\}$ is known, so is J_i and W_{ij}. Moreover, τ_{ijv} is observed, while λ^f_i is recovered from the matching function (λ^f_i is the average ratio, $\frac{1}{T}\sum_t m_{it}/f_{it}$, where f_{it} is estimated). We calibrate the freight survival probability

[32] We assume that our data comes from one steady state, so that (s_t, f_t, s^w_t) is fixed at (s^*, f^*, s^{w*}); hence $P_{ij}(\theta)$ does not depend on t (it also does not depend on the ship l since ships are homogeneous). During our sample period of 2012-2016 the industry did not experience any major shocks. We also estimated the model separately by season to allow for seasonal time variation and find our results to be similar across the four seasons.

to $\delta = 0.99$. Now, equation (2.26) has two unknowns: the valuation v and the bargaining coefficient γ. First, we pin down γ from external information on the average value of international trade in commodities and obtain $\gamma = 0.3$.[33] Now, given this estimate for γ, we recover exporter valuations point-wise from (2.26) and can obtain their distribution, F_{ij}^v, nonparametrically.[34] Note that valuations are drawn from an origin-destination specific distribution, which allows for arbitrary correlation between a cargo's valuation and destination.

The exporter costs κ_{ij} capture both the cost of production, as well as any export costs beyond shipping prices and are estimated from the exporters' chosen destinations. Indeed, given the choice probabilities $\tilde{G}_{ij}$ defined in (2.12) we can recover κ_{ij} as follows (Berry (1994)):

$$\kappa_{ij} = J_{ij}^f - \left(\ln \tilde{G}_{ij} - \ln \tilde{G}_{i0}\right) \tag{2.27}$$

where J_{ij}^f is now known; indeed, recall from (2.20) that $J_{ij}^f = \alpha_i(\mu_{ij} - \tau_{ij})$ and we can now calculate $\alpha_i = \lambda_i^f/(1 - \beta\delta(1 - \lambda_i^f))$, as well as the mean valuations μ_{ij}. Finally, the satellite data provides direct information on the frequencies G_{ij}, the proportion of loaded trips from i to j. We do, however, need to determine the share of the outside option or equivalently, the number of entrants d_i and potential entrants $\mathcal{E}_i$, in order to compute $\tilde{G}_{ij} = G_{ij}(1 - \tilde{G}_{i0})$. We obtain the number of entrants by solving for

[33] To obtain the average valuation worldwide, we first collect the average price of the five most common commodities (iron ore, coal, grain, steel and urea) from Index Mundi, and multiply it by the average tonnage carried by a bulk carrier (this is equal to the average bulker size times its utilization rate; see Footnote 35). We then set $\bar{\mu}$ as their weighted average based on each commodity's frequency in shipping contracts; we find $\overline{\mu}$ to equal 7 million US dollars. Finally, solving (2.26) with respect to γ and averaging over i, j, v yields:

$$\gamma = \frac{(1 - \beta\delta)\,(\overline{\mu} - \overline{\tau})}{\beta\delta E_{ijv}\lambda_i^f \tau_{ijv} + (1 - \beta\delta)\,\overline{\mu} - E_{ij}\left(1 - \beta\delta\left(1 - \lambda_i^f\right)\right)(J_i - W_{ij})}$$

where $\bar{\tau}$ is the average observed price.

[34] We have assumed that the exporter obtains zero payoff when he does not find a match. It is not possible to separately identify valuations from such inventory costs or scrap values.

d_i from the freight transition (2.1) and taking the average. The number of potential entrants $\mathcal{E}_i$ is set equal to the total production of the relevant commodities for each region i.[35]

2.6 Results

In this section we present the results from our empirical analysis. Throughout the estimation, we consider 15 geographical regions, depicted in Figure 2.15 in the Appendix.[36] To determine the regions, we employ a clustering algorithm that minimizes the within-region distance between ports.

2.6.1 Matching Function

Search Frictions Test Before presenting the main results, we provide a simple test for search frictions, inspired by the model and the empirical methodology outlined in Section 2.5.1. Suppose that for some region i, it is known that there are more ships than exporters, i.e. $\min(s_{it}, f_{it}) = f_{it}$. If there are no search frictions, so that $m_{it} = \min(s_{it}, f_{it}) = f_{it}$, exogenously changing the number of ships does not affect the number of matches. In contrast, if there are search frictions, any exogenous change in the number of ships changes the number of matches. We can thus test for search frictions by using the exogenous changes in ocean weather conditions to explore whether changing the number of ships in regions with a lot more ships than

[35] We collect annual country-level production data for grain (FAO), coal (EIA), iron ore (US Geological Survey), fertilizer (FAO) and steel (World Steel Association). To transform the production tons into a number of potential freights (i.e. shipments that fit in our bulk vessels), we first scale the production to adjust for the coverage of our data (we observe about half of the total fleet) and then divide by the average "active" ship size. A ship operates on average 340 days per year (due to maintenance, repairs, etc.) and has a deadweight utilization of about 65%. A region's total production serves as an upper bound to the region's exports.

[36] The trade-off here is that we need a large number of observations per region, while allowing for sufficient geographical detail. The regions are: West Coast of North America, East Coast of North America, Central America, West Coast of South America, East Coast of South America, West Africa, Mediterranean, North Europe, South Africa, Middle East, India, Southeast Asia, China, Australia, Japan-Korea. We ignore inter-regional trips.

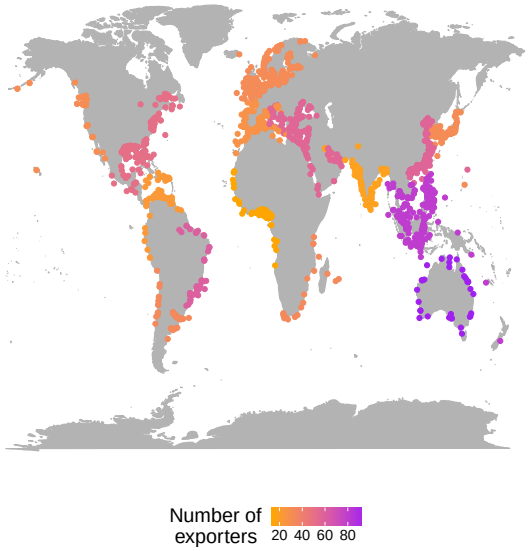

Figure 2.5: Average weekly number of exporters.

exporters affects the realized number of matches. To proxy for weather conditions, we employ the unpredictable component of wind at sea.[37] Since we do not observe freights directly, for each region we consider periods when there are substantially more ships than matches (though recall that overall during our sample period there is substantial ship excess supply). Table 2.4 in the Appendix presents the results across regions during weeks with at least twice as many ships as matches. We find that indeed matches are affected by weather conditions, consistent with the presence of search frictions.

Matching Function Estimates We now turn to the results from our main methodology of Section 2.5.1. Table 2.5 in the Appendix presents the results from the first stage regression of the number of ships on the weather for all regions and reveals that ocean wind has significant impact. Figure 2.5 presents the weekly average

[37]In particular, we divide the sea surrounding each region into 8 different zones. For each zone we use information on the wind speed at different distances from the coast and in different directions. To obtain the unpredictable component of weather we run a VAR regression of these weather indicators on their lag component and season fixed effects. We experiment with the lag structure and the results are overall robust. Finally, the results are robust to running the VAR jointly for neighboring zones.

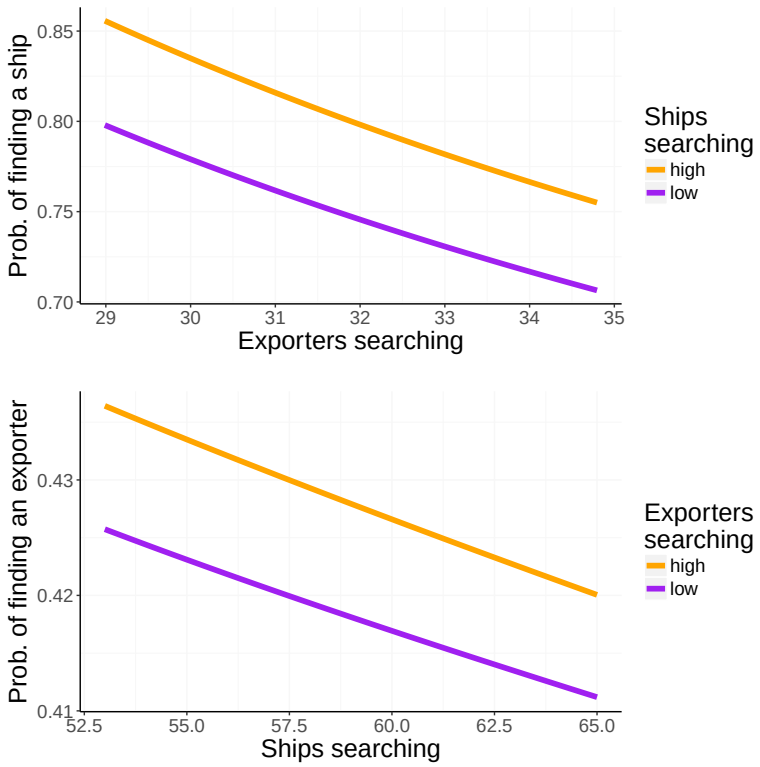

Figure 2.6: Estimated match finding probabilities for ships and exporters, λ and λ^f in the West Coast of South America. "High" corresponds to the 40th percentile, while "low" corresponds to the 60th percentile.

number of exporters in the world. Exporters are concentrated in Australia, the East Coast of South America and Southeast Asia. India, Africa and Central America have the fewest freights.

To visualize the matching function, Figure 2.6 plots matching rates for ships and exporters, for the West Coast of South America as one example. The top panel plots the matching rate for exporters, λ^f, as a function of the number of exporters searching and for different levels of ships. Note that, as expected, λ^f declines as the market gets crowded with exporters. Similarly, the bottom panel plots the probability that a ship finds a match, λ, as a function of the number of ships and for different levels of exporters. Again, this probability declines in the number of searching ships. It is also worth noting, that exporters have substantially higher chances of finding a match than ships, consistent with our sample period of high ship supply and low demand,

as well as our conservative scale restriction on the exporters (see footnote 30). This is true in all regions.

To measure the extent of search frictions in different regions, we compute the average percentage of weekly "unrealized" matches; i.e. $(\min\{s_i, f_i\} - m_i) / \min\{s_i, f_i\}$. The results are plotted in Figure 2.7 and reveal that search frictions are heterogeneous over space and may be sizable, with up to 20% of potential matches "unrealized" weekly in regions like West Africa and parts of South America and Europe. On average, 17.2% of potential matches are "unrealized".[38,39] Figure 2.16 in the Appendix, depicts these ratios for all regions, as well as the minimum between ships and freights, along with confidence intervals constructed from 500 bootstrap samples. We can reject that the matching function equals the minimum for all regions.

The ratio of "unrealized" matches correlates well with the observed within-region price dispersion, an indicator of search frictions. It also correlates with the ratio of incoming and outgoing ships, discussed in Section 2.3.3, as for instance there are more "unrealized" matches in Chile than Norway. When we estimate the matching function separately for Capesize (biggest size) and Handysize (smallest size) vessels, we find that for Capesize, where the market is thinner, the ratios of "unrealized" matches are lower.

2.6.2 Ship Costs

In our baseline specification, we construct seven groups for the sailing cost c^s_{ij}, roughly based on the continent and coast of the origin; and we estimate all port wait costs

[38]Results are overall robust if instead of imposing that the matching function is homogeneous of degree one, we fix the distribution of f; see Footnote 29. In our case, a $[0, 1]$ uniform distribution for freights does not sound plausible since we need to also satisfy $m_{it} \leq f_{it}$ for all i, t. Therefore, we instead experimented with more flexible distributions (normal, log-normal) and calibrated their parameters so that this inequality is always just satisfied- this again yields the most conservatively estimated level of search frictions.

[39]It is worth noting that this does not imply that in the absence of search frictions we would have 17.2% more matches; this is simply a measure of the severity of search frictions in different regions.

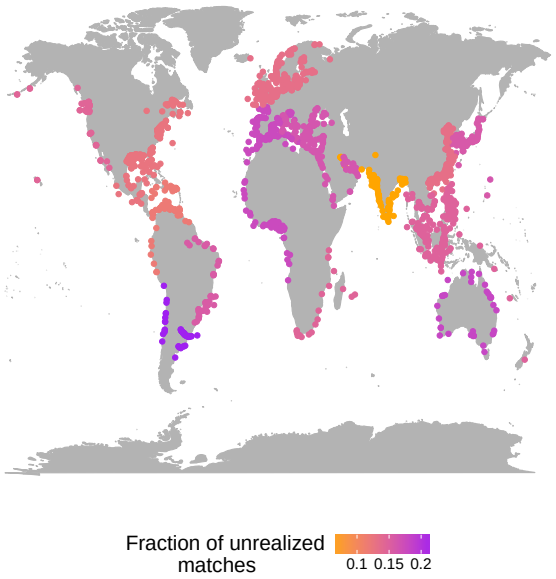

Figure 2.7: Average weekly share of "unrealized" matches because of search frictions.

c_i^u, for all i.[40] Note that c_{ij}^s is the *per week* sailing cost from i to j and its major component is the cost of fuel. We thus set this cost for one of the groups (for trips originating from the East Coast of North and South America) equal to the average fuel price.[41] Note also that since the fuel cost is paid by the exporter when the ship is loaded, we add it to the observed prices.

The first two columns of Table 3.1 report the results.[42] Not surprisingly, sailing costs are fairly homogeneous. Port wait costs are more heterogeneous and large, ranging between 130,000 and 260,000 US dollars per week. Consistent with industry narratives, it is costly to let a ship waiting at port, both due to direct port and security

[40]The seven groups are: (i) Central America, West Coast Americas; (ii) East Coast Americas; (iii) West and South Africa; (iv) Mediterranean, Middle East and North Europe; (v) India; (vi) Australia and Southeast Asia; (vii) China, Japan and Korea.

[41]The average weekly price of fuel is 69,100$. We have experimented heavily with different types of identification restrictions and the results are robust. In particular, we have considered (i) $c_{ij}^s = c^s$ for all i and j; (ii) coarser and finer groups; (iii) c_{ij}^s clustered by the distance between i and j, as well as clustered by the weather between i and j (both capture nonlinear effects of distance on the sailing cost). The ship also incurs operating costs (crew, maintenance, etc.). However, these are fixed costs of operation; as such they do not affect the ships' decisions and can be ignored.

[42]The standard errors are computed from 500 bootstrap samples with the resampling done at the ship level. We combine these bootstrap samples with those of the matching function to incorporate the error from the matching function estimation.

fees, as well as the rapid depreciation of the ship's machinery and electronics. Ports in the Americas are the most expensive, while ports in China, India, Southeast Asia and the Middle East are the cheapest. The standard deviation of the preference shocks, σ, is estimated at about 11,000 US dollars, roughly 10% of price, which implies that the preference shocks do not account for a disproportionately large part of utility or ballast decisions. As shown in Figure 2.17 of the Appendix, the fit is very good, as our predicted choice probabilities are very close to the observed ones.[43]

2.6.3 Exporter Valuations and Costs

In Figure 2.8 we plot the average exporter valuations across origins, while the third column of Table 3.1 reports the estimates. There is substantial heterogeneity in valuations across space. South and North America have the highest valuations, while Europe and Southeast Asia have the lowest. This ranking is reasonable, as for instance, Brazil exports grain which is expensive, while Southeast Asia exports mostly coal, which is one of the cheapest commodities. We generalize this example by focusing on grain, the most expensive frequently shipped commodity. In particular, using data from Comtrade, we explore whether countries that have a high share of grain exports tend to have higher estimated valuations. The results, shown in Figure 2.9, reveal that indeed there is a positive correlation between the two, suggesting that exporters with higher valuations may be producers of more expensive products.[44] Of course, there may be other factors determining the valuation of an exporter such as inventory control, just in time production, etc. On average, the average price τ_{ij} is

[43] As our data comes from a period of historically low shipping prices, our estimated value functions are negative. This is partly due to the fact that we are not modeling ships' expectations so shipowners do not realize that under a mean-reverting demand for seaborne trade prices will go up eventually (see Kalouptsidi (2014)). If we compute the equilibrium under higher exporter valuations that lead to prices closer to the ones observed before 2010, the ship value function indeed becomes positive.

[44] As discussed in Section 2.3.3, the dependence of prices on the type of good is suggestive of search frictions in the market. In our estimation we are able to back out the valuations given our estimates for search frictions. The external validation discussed in this paragraph also supports our conclusions that search frictions are present in these markets.

equal to about 5% of the mean valuation μ_{ij}, consistent with other estimates in the literature (e.g. UNCTAD (2015), Hummels et al. (2009)).

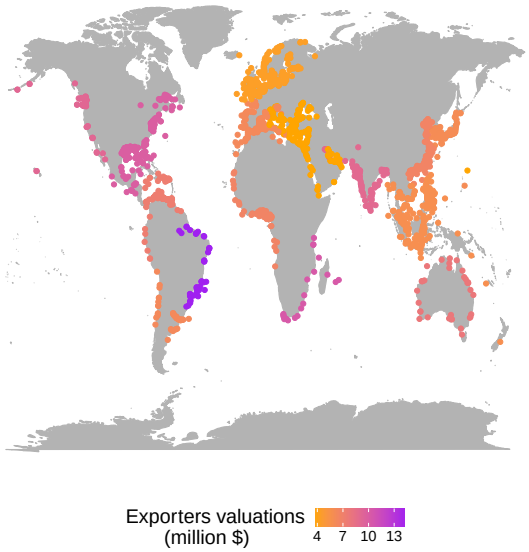

Figure 2.8: Average exporter valuations.

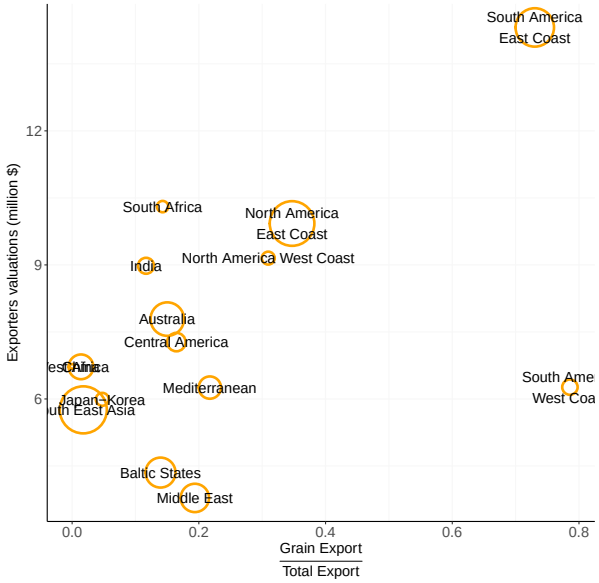

Figure 2.9: Average exporter valuations and share of exports in grain (source, Comtrade). The size of the circle proxies the number of observations.

Finally, we turn to exporter costs. The share of the exporters' outside option, computed from the total commodity production in the region, is on average 62%.

China and India (South America, Australia and Southeast Asia) feature the highest (lowest) outside share, consistent with their high (low) imports. The estimated exporter costs exhibit substantial heterogeneity across destinations from a given origin, as well as across origins. On average κ_{ij} is the same order of magnitude as the average valuation μ_{ij}. Moreover, we find that exporter costs are lower between an origin i and a destination j if the same language is spoken at i and j, which is reasonable since κ_{ij} includes both production costs, as well as other exporting costs.

Our estimation results are robust when we estimate the costs separately by season. Moreover, we have performed our estimation and counterfactuals for Handymax and Handysize vessels alone (the segments for which we have sufficient data at the baseline region-week level) and found that the resulting counterfactuals are very similar.

	Port Costs c_u	Sailing Costs c_s	Exporters Valuations μ_v	Preference Shock σ
North America West Coast	2.458	0.693	79.605	
	(0.07)	(0.002)	(2.038)	
North America East Coast	2.271	-	103.145	
	(0.021)	-	(2.229)	
Central America	1.846	0.693	73.161	
	(0.022)	(0.002)	(3.007)	
South America West Coast	1.996	0.693	59.063	
	(0.026)	(0.002)	(1.679)	
South America East Coast	2.563	-	125.877	
	(0.027)	-	(3.001)	
West Africa	1.421	0.64	57.838	
	(0.015)	(0.002)	(2.658)	
Mediterranean	1.637	0.568	59.87	
	(0.018)	(0.003)	(2.475)	
North Europe	1.399	0.568	49.199	
	(0.009)	(0.003)	(1.959)	
South Africa	2.478	0.64	99.074	
	(0.035)	(0.002)	(2.907)	
Middle East	1.273	0.568	44.203	
	(0.007)	(0.003)	(2.355)	
India	1.48	0.624	84.722	
	(0.014)	(0.003)	(4.2)	
South East Asia	1.67	0.56	72.282	
	(0.008)	(0.002)	(3.324)	
China	1.438	0.558	66.382	
	(0.01)	(0.002)	(3.61)	
Australia	2.635	0.56	70.507	
	(0.025)	(0.002)	(2.543)	
Japan-Korea	1.53	0.558	55.589	
	(0.022)	(0.002)	(2.514)	
				0.117
				(0.0008)

Note: all the estimates are in 100,000 USD. Standard errors computed from 500 bootstrap samples.

Table 2.3: Ship costs and exporter valuation estimates. The sailing cost for the East Coast of North and South America is set equal to 0.69 (the fuel cost).

2.7 Counterfactuals

In this section, we use our estimated model to explore a number of questions of interest through the lens of endogenous trade costs. We consider: the change in exports following an improvement in shipping efficiency, a Chinese slow-down, the opening of the Northwest Passage and the trade reduction due to search frictions.

These counterfactuals illustrate three novel mechanisms of our setup. First, as model primitives change and directly affect the value of a match, the ship's outside option, J_i, is also affected and provides a new channel that impacts trade costs and exporting. Second, the economic and geographic network of countries matters. For instance, when a shock makes a trading partner j more attractive for ships, say because it offers more options to ships that end their trip there, then W_{ij} increases and from equation (2.17), prices to j go down and exports increase. Similarly, exporting countries that are close to large net importers, benefit from the "glut" of ships that end their trip there and wish to ballast elsewhere. Third, how each country is affected in every one of the counterfactuals below depends crucially on (i) the size of its trade imbalance and (ii) its mean value, v, especially when it is a net exporter. For instance, as we will see below, under improvements in shipping technology or matching efficiency, ship reallocation leads to "polarization", whereby more ships are now more likely to reallocate to large, high value exporting countries

To perform the counterfactuals, we compute the steady state spatial equilibrium distribution of ships and exporters. In the Appendix, we provide the computational algorithm employed.[45]

2.7.1 Change in Trade due to a Change in Shipping Efficiency

We consider how a 10% decrease in the cost of shipping, c^s, affects shipping prices, τ, and trade flows. This change in transportation costs has two effects:

[45] Our algorithm always converged to the same solution, even from very different starting values.

First, there is a direct increase in the surplus of all matches, since now a match between a ship and a freight is more valuable.[46] All else equal, this reduces export prices, τ, which in turn increases the value of an unmatched exporter, J^f, and thus induces more entry into the export market.

Second, reducing c^s implies that ballasting is cheaper and ships can reallocate across space more freely. Therefore, their "outside option", J, is now higher. Since Nash bargaining requires that both parties receive their outside option plus a share of the surplus, an increase in the ship's outside option leads, all else equal, to an increase in prices, τ (see the price equation (2.17)). Put differently, reduced transportation costs imply that ships are less "tied" to their current region and as freights' monopsony power is reduced, ships receive higher prices. This effect tends to mitigate the increase in freight entry driven by the direct effect on the surplus.

Figure 2.10 presents the results and showcases that there is considerable heterogeneity in different regions' reaction to a change in transportation costs. The impact of the second effect on prices is easily seen, as prices decline by at most 3%, instead of 10% (the decline in c^s). North America, the East Coast of South America, Australia and Southeast Asia see the highest increase in exporting (5-7%). India, China, Japan and Southern Europe are among the regions where exporting increases the least (1-3%). We find that the second effect, capturing the change in ships' outside option, correlates both with a country's net exporting status, as well as the average valuation of its exporters. Indeed, net exporters experience a higher increase in their exports, than net importers all else equal. Ships ending a trip at net importing countries now face lower ballast costs and are thus less likely to wait there for a cargo; their outside option is higher and they can command higher prices. For

[46] Formally, using the ship and freight value functions, the match surplus is given by

$$S_{ijv} = v - \frac{c^s_{ij}}{1-\beta(1-\xi_{i,j})} + \frac{\xi_{ij}\beta}{1-\beta(1-\xi_{ij})} U_j - J_i - J^f_{ijv}.$$

A decline in c^s_{ij}, holding everything else constant, directly increases S_{ijv}.

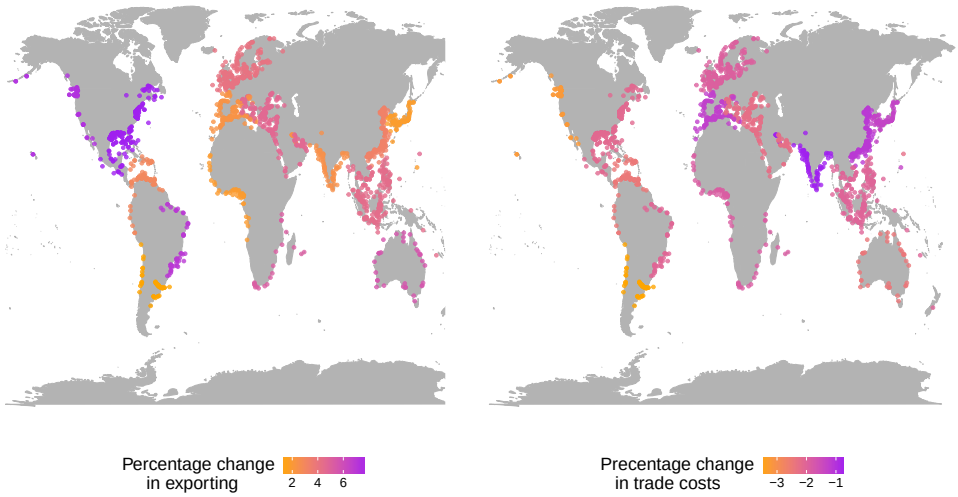

Figure 2.10: Change in exports and trade costs under a 10% decline in the sailing cost c^s.

instance, we find that in India, China and Japan the second effect mutes the direct effect substantially: if ships could not reallocate, the increase in exports in these countries would have been three times higher. In contrast, high value net exporters (e.g. Northeast America, as well as Brazil) benefit from the increased willingness of ships to ballast, and they do see an increase in the number of ships ballasting there: as transport costs, and thus distance, now matter less, mean freight valuations, μ, become a relatively more important determinant of ships' decisions.

2.7.2 Chinese Slow-down

We next explore how shocks propagate in a world where trade costs are endogenous by considering a Chinese slow-down (i.e. a reduction in the mean valuation of freights going to China, $\mu_{i,china}$, by 10%). Figure 2.11 presents the resulting trade costs and trade flows across the world.

We begin discussing the results by looking at China itself: shipping prices from China increase by 2%, while exporting declines by 6%. This illustrates the complementarity between imports and exports: the high Chinese imports, led to a large

number of ships ending their trip in China and looking for a freight there, which led to reduced trade costs for Chinese exporters. Therefore, when imports decline, fewer ships end up in China and Chinese exporters are hurt.

Next, note that China's trading partners, such as Australia, Southeast Asia and Brazil, naturally experience a substantial decline in their exports (16-40%). In addition, much like in the previous counterfactual, the decline in exporting is dampened by the reduction in ships' outside option as ships are overall worse off. Prices also decline, both because China is a relatively expensive destination ($\mu_{i,China}$ is high), but also because ships' outside option is lower.

Finally, it is instructive to trace out the geographically heterogeneous response of exports to this shock, in order to investigate the role of the network of countries. The dampening accruing from ships' lower outside options is much lower for the big exporters close to China, such as Australia and Southeast Asia, and much larger for further away exporters, such as Brazil. This underscores the importance of being close to a large net importer like China: exporting countries in that "pocket" of the world, gain not just by directly exporting to China, but also indirectly from the increased supply of ships in that region. Indeed, the Southeast Pacific region consists of some close-by large exporters (Australia, Southeast Asia) and importers (China, Japan, India) and thus benefits from a "cheap" supply of ships that remains in the area ballasting and trading between these countries. When Chinese demand falls these ships reallocate to other parts of the world and tend to push up exports there, dampening the overall decline. To see this, we compute that if ships were not able to reallocate and ships' outside options had not changed, Brazilian exports would have fallen by 21% rather than 15%, while in Northeast America exports would have declined by 15% rather than 9%. In contrast, Chinese exports would have barely been affected, while the decline in Southeast Asian, Japanese and Australian exports would have been roughly the same.

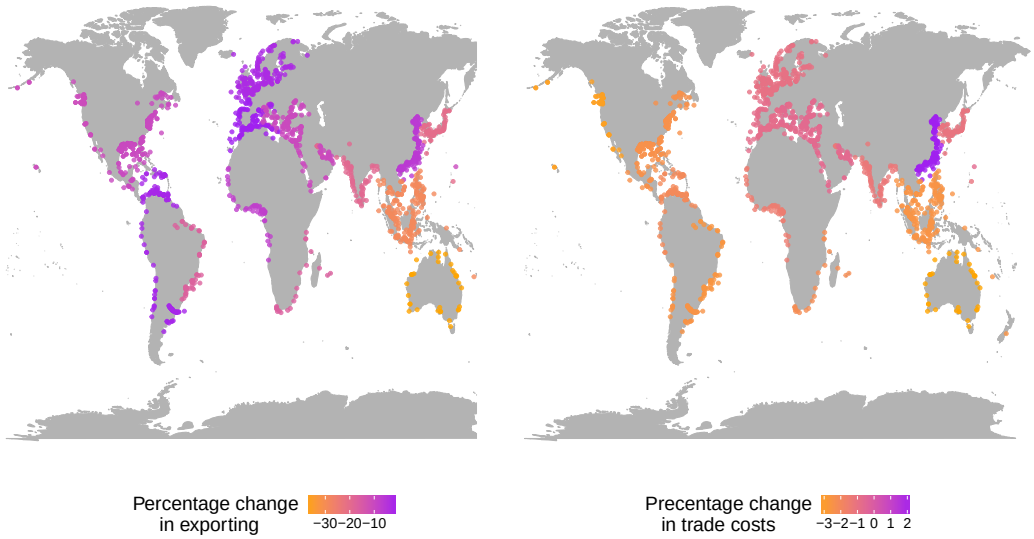

Figure 2.11: Change in exports and trade costs under a Chinese slow-down.

2.7.3 Opening the Northwest Passage

The Northwest Passage is a sea route connecting the northern Atlantic and Pacific Oceans through the Arctic Ocean, along the northern coast of North America. This route is not easily navigable due to Arctic sea ice; with global warming and ice thinning, there is public discussion about opening the passage to be exploited for shipping. The Northwest Passage would reduce the travel costs between Northeast America and the Far East.

To simulate the opening of the Northwest Passage, we reduce transport costs between the East Coast of North America and China/Japan/S.Korea by 30%.[47] Figure 2.12 presents the resulting change in exports by region. Naturally, Northeast America sees its exporting rise, by 8.6%. Interestingly, exports from China and Japan/S.Korea fall, as ballasting is now less costly for ships: when in China or Japan they can now ballast to Northeast America more cheaply. Indeed, we find that the probability of an unmatched ship staying in China (Japan/S. Korea), P_{ii}, is now 25%

[47] We calculate the change in sailing costs of traveling via the Northwest Passage from Ostreng et al. (2013).

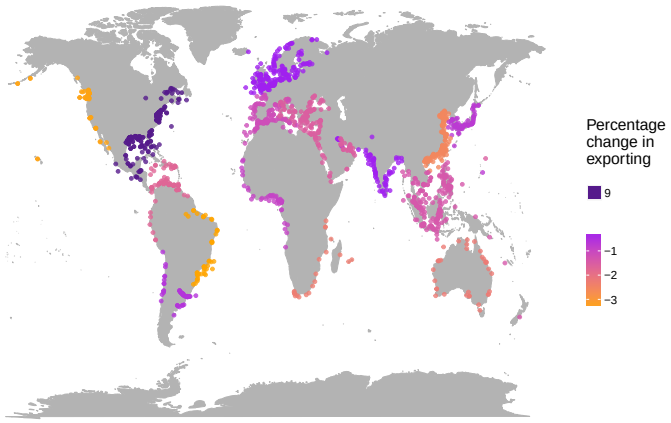

Figure 2.12: Change in exports if the Northwest Passage becomes a viable commercial route.

(28%) lower. The ships' higher outside option, tends to increase prices and decrease exporting.

Finally, Figure 2.12 reveals that other countries, not directly affected by the opening of the Northwest Passage experience changes in their trade. Overall, exports decline by up to about 3%. This decline is due to the ships' higher outside option. The response of different countries is quite heterogeneous and illustrates the neighbor and network effects. For instance, we see that Brazil, as well as Northwest America are hurt the most. Indeed, as these countries are close to Northeast America ships that used to ballast there to load now prefer to ballast to Northeast America instead. We find that the number of ships that ballast to Brazil (Northwest America) is lower by about 2% (3%), while ballasting to Northeast America is 13% higher. In contrast, Australian and Southeast Asian exports do not decline as much, as these countries are shielded by their closeness to China and India. These effects are present because of the endogenous trade cost and demonstrate the spillover of such policy changes through the network of countries.

2.7.4 The Trade Lost because of Search Frictions

Finally, we quantify the trade lost because of search frictions. To do so, we shut down search frictions by setting $m(s, f) = \min\{s, f\}$.[48] Figure 2.13 presents the resulting change in exports by region, as well as the change in ballasting. We find that exporting would be 6-45% higher across different regions in the world with an average increase of 23%.[49] It is of course not surprising that we see higher exports globally in the absence of search frictions. Interestingly, the change in trade is disperse geographically. While countries that experience more severe frictions, as captured in Figure 2.7, roughly experience somewhat larger increases in exports, this is not always the case. Indeed, a country's net exporting status is a more important determinant of the extent to which the country can attract economic activity. Large net exporters like Brazil and the Northeast America experience disproportionally large increases in exports, as differences in frictions across regions are no longer relevant and exporting size becomes a more important determinant of trade. The correlation between the change in exports and a country's trade imbalance is 0.9. Since search frictions have been removed, ships ballast more towards large net exporters. Indeed, in a world with no search frictions, as shown in in Figure 2.13, large net exporters such as Brazil, Australia and North America experience large increases in the number of ships ballasting there. As previously discussed, when impediments to trade are reduced, country differences in their export status and size become relatively more important determinants of ships' ballast decisions and the resulting trade flows.

[48] In practice we set $m(s, f) = \alpha \min\{s, f\}$ with $\alpha = 0.99$, in order to maintain the Nash bargaining setup and the model comparable.

[49] It is worth noting that the impact of search frictions is quantitatively important also because a ship completes several trips during the course of a year and therefore finds itself "unemployed" extremely often; in contrast, in labor markets, the average worker experiences unemployment once every few years.

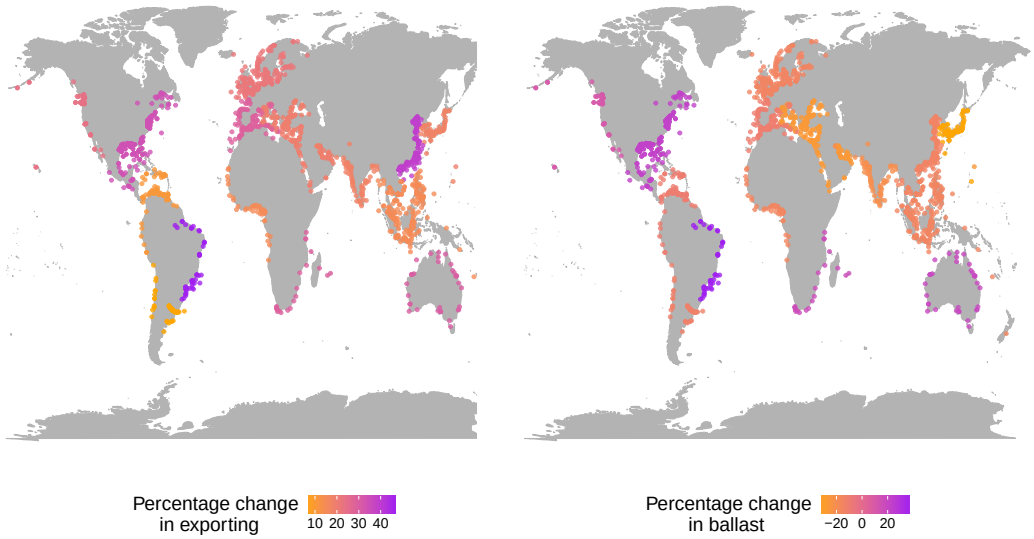

Figure 2.13: Change in exporting in a world without search frictions.

2.8 Conclusion

In this paper, we build a dynamic spatial search model for world ships and exporters. Using unique data on shipping contracts and ship movements we recover the main primitives of interest: the matching function between ships and freights, the distribution of searching exporters, ship costs and exporter valuations. Our methodology allows us to obtain the matching process flexibly, without relying on assumptions regarding the extent of search frictions or the parametric form of the matching function. We demonstrate that accounting for the endogeneity of trade costs is important in both descriptive analysis (e.g. elasticities, shock propagation), as well as policy analysis (e.g. transportation infrastructure planning). Finally, we find that search frictions substantially reduce world trade.

References

Aguirregabiria, V. and Mira, P.: 2007, Sequential estimation of dynamic discrete games, *Econometrica* **75**(1), 1–53.

Allen, T. and Arkolakis, C.: 2014, Trade and the topography of the spatial economy, *Quarterly Journal of Economics* **129**(3), 1085–1140.

Allen, T. and Arkolakis, C.: 2016, The welfare effects of transportation infrastructure improvements, *mimeo, Yale University* .

Anderson, J. E. and Van Wincoop, E.: 2003, Gravity with gravitas: A solution to the border puzzle, *American Economic Review* **93**(1), 170–192.

Asturias, J.: 2016, Endogenous transportation costs, *mimeo, School of Foreign Service in Qatar, Georgetown University* .

Bajari, P. and Benkard, C. L.: 2005, Demand estimation with heterogeneous consumers and unobserved product characteristics: A hedonic approach, *Journal of Political Economy* **113**(6), 1239–1276.

Bajari, P., Benkard, C. L. and Levin, J.: 2007, Estimating dynamic models of imperfect competition, *Econometrica* **75**(5), 1331–1370.

Berry, S.: 1994, Estimating discrete-choice models of product differentiation, *RAND Journal of Economics* **25**(2), 242–262.

Bertsekas, D. P. and Tsitsiklis, J. N.: 2015, *Parallel and Distributed Computation: Numerical Methods*, Athena Scientific.

Buchholz, N.: 2016, Spatial equilibrium, search frictions and efficient regulation in the taxi industry, *mimeo, Princeton University* .

Burdett, K. and Mortensen, D. T.: 1998, Wage differentials, employer size, and unemployment, *International Economic Review* **39**(2), 257–273.

Collard-Wexler, A.: 2013, Demand fluctuations in the ready-mix concrete industry, *Econometrica* **81**(3), 1003–1037.

Donaldson, D.: 2012, Railroads of the raj: Estimating the impact of transportation infrastructure, *forthcoming, American Economic Review* .

Donaldson, D. and Hornbeck, R.: 2016, Railroads and american economic growth: A "market access" approach, *Quarterly Journal of Economics* **131**(2), 799–858.

Eaton, J., Jinkins, D., Tybout, J. and Xu, D.: 2016, Two-sided search in international markets, *mimeo, Penn State University* .

Eaton, J. and Kortum, S.: 2002, Technology, geography and trade, *Econometrica* **70**(5), 1741–1779.

Ericson, R. and Pakes, A.: 1995, Markov-perfect industry dynamics: A framework for empirical work, *The Review of Economic Studies* **62**(1), 53–82.

Fajgelbaum, P. D. and Schaal, E.: 2017, Optimal transport networks in spatial equilibrium, *NBER working paper* **23200**.

Frechette, G. R., Lizzeri, A. and Salz, T.: 2016, Frictions in a competitive, regulated market evidence from taxis, *mimeo, New York University* .

Head, K. and Mayer, T.: 2004, The empirics of agglomeration and trade, *in* J. V. Henderson and J. F. Thisse (eds), *Handbook of Regional and Urban Economics*, Vol. 4, Elsevier.

Hopenhayn, H. A.: 1992, Entry, exit, and firm dynamics in long run equilibrium, *Econometrica* **60**(5), 1127–1150.

Hosios, A. J.: 1990, On the efficiency of matching and related models of search and unemployment, *The Review of Economic Studies* **57**(2), 279–298.

Hotz, V. J. and Miller, R. A.: 1993, Conditional choice probabilities and the estimation of dynamic models, *Review of Economic Studies* **60**(3), 497–529.

Hummels, D.: 2007, Transportation costs and international trade in the second era of globalization, *The Journal of Economic Perspectives* **21**(3), 131–154.

Hummels, D., Lugovskyy, V. and Skiba, S.: 2009, The trade reducing effects of market power in international shipping, *Journal of Development Economics* **89**(1), 84–97.

Hummels, D. and Skiba, S.: 2004, Shipping the good apples out? an empirical confirmation of the alchian-allen conjecture, *Journal of Political Economy* **112**(6), 1384–1402.

Ishikawa, J. and Tarui, N.: 2015, Backfiring with backhaul problems: Trade and industrial policies with endogenous transport costs, *mimeo, University of Hawaii* .

Kalouptsidi, M.: 2014, Time to build and fluctuations in bulk shipping, *American Economic Review* **104**(2), 564–608.

Kalouptsidi, M.: 2017, Detection and impact of industrial subsidies, the case of chinese shipbuilding, *forthcoming, Review of Economic Studies* .

Kalouptsidi, M., Scott, P. T. and Souza-Rodrigues, E.: 2016, Identification of counterfactuals in dynamic discrete choice models, *NBER working paper* **21527**.

Krugman, P.: 1991, Increasing returns and economic geography, *The Journal of Political Economy* **99**(3), 483–499.

Lagos, R.: 2000, An alternative approach to search frictions, *The Journal of Political Economy* **108**(5), 851–873.

Lagos, R.: 2003, An analysis of the market for taxicab rides in new york city, *International Economic Review* **44**(2), 423–434.

Lange, F. and Papageorgiou, T.: 2017, Estimating matching functions when search intensity is unobserved, *mimeo, McGill University* .

Matzkin, R. L.: 2003, Nonparametric estimation of nonadditive random functions, *Econometrica* **71**(5), 1339–1375.

Mortensen, D. T.: 2003, *Wage Dispersion: Why Are Similar Workers Paid Differently?*, MIT Press, Cambridge, MA.

Mortensen, D. T. and Pissarides, C. A.: 1994, Job creation and job destruction in the theory of unemployment, *The Review of Economic Studies* **61**(3), 397–415.

Ostreng, W., Eger, K. M., Floistad, B., Jorgensen-Dahl, A., Lothe, L., Mejlaender-Larsen, M. and Wergeland, T.: 2013, *Shipping in Arctic Waters: A Comparison of the Northeast, Northwest and Trans Polar Passages*, Springer, Berlin, Heidelberg.

Pakes, A., Ostrovsky, M. and Berry, S. T.: 2007, Simple estimators for the parameters of discrete dynamics games, *The RAND Journal of Economics* **38**(2), 373–399.

Petrongolo, B. and Pissarides, C. A.: 2001, Looking into the black box: A survey of the matching function, *Journal of Economic Literature* **39**(2), 390–431.

Pissarides, C. A.: 2000, *Equilibrium Unemployment Theory*, 2nd edn, MIT Press, Cambridge, MA.

Postel-Vinay, F. and Robin, J.-M.: 2002, Equilibrium wage dispersion with worker and employer heterogeneity, *Econometrica* **70**(6), 2295–2350.

Reyes-Heroles, R.: 2016, The role of trade costs in the surge of trade imbalances, *mimeo, Federal Reserve Board* .

Rogerson, R., Shimer, R. and Wright, R.: 2005, Search-theoretic models of the labor market: A survey, *Journal of Economic Literature* **43**(4), 959–988.

Rust, J.: 1987, Optimal replacement of gmc bus engines: An empirical model of harold zurcher, *Econometrica* **55**(5), 999–1033.

Ryan, S. P.: 2012, The costs of environmental regulation in a concentrated industry, *Econometrica* **80**(3), 1019–1062.

UNCTAD: 2015, Review of maritime transport, *Technical report*, United Nations Conference on Trade and Development, Geneva, Switzerland.

Waugh, M.: 2010, International trade and income differences, *American Economic Review* **100**(5), 2093–2124.

Wong, W. F.: 2017, The round trip effect: Endogenous transport costs and international trade, *mimeo,University of Wisconsin-Madison* .

WTO: 2015, International trade statistics, *Technical report*, World Trade Organization, Geneva, Switzerland.

Appendix

2.A Construction of Ship Travel Histories

Here, we describe the construction of ships' travel histories. The first task is to identify stops that ships make, using the EE data which provides the exact location of ships every six minutes. A stop is defined as an interval of at least 24 hours, during which (i) the average speed of the ship is below 5 mph (the sailing speed is between 15 and 20 mph) and, (ii) the ship is located within 250 miles from the coast. A trip is the travel between two stops.

The second task is to identify whether a trip is loaded or ballast. To do so, we use the ship's draft: high draft indicates that a larger portion of the hull is submerged and therefore the ship is loaded. The distribution of draft for a given vessel is roughly bimodal, since as described in Section 2.2, a hired ship is usually fully loaded. Therefore, we can assign a "high" and a "low" draft level for each ship and consider a trip loaded if the draft is high (in practice, the low draft is equal to 70% of the high draft). As not all satellite signals contain the draft information, we consider a trip ballast (loaded) if we observe a signal of low (high) draft during the period that the ship is sailing. If we have no draft information during the sailing time, we consider the draft at adjacent stops. Finally, we exclude stops longer than six weeks, as such stops may be related to maintenance or repairs.

The third and final task is to refine the origin and destination information provided in the Clarksons contracts. Although the majority of Clarksons contracts provide some information on the origin and destination of the trip, this is often vague (e.g. "Far East", "Japan-S. Korea-Singapore"), especially in the destinations. We use the EE data to refine the contracted trips' origins and destinations by matching each Clarksons contract to the identified stop in EE that is closest in time and, when possible, location. In particular, we use the loading date annotated on each contract to find a stop in the ship's movement history that corresponds to the beginning of the contract. For destinations where our information for Clarksons is noisy we search the ship's history for a stop that we can classify as the end of the contract. In particular, we consider all stops within a three month window (duration of the longest trip) since the beginning of the contract. Among these stops we eliminate all those that (i) are in the same country in which the ship loaded the cargo and (ii) are in Panama, South Africa, Gibraltar or at the Suez canal and in which the draft of arrival is the same as the draft of departure (to exclude cases in which the ship is waiting to pass through a strait or a canal). To select the end of the contract among the remaining options we consider the following possibilities:

1. If the contract reports a destination country and if there are stops in this country, select the first of these stops as the end of the trip;

2. If the destination country is "Japan-SKorea-Singapore", and if there are stops in either Japan, China, Korea, Taiwan or Singapore, we select the first among these as the end of the trip;

3. If the contract does not report a destination country and there are stops in which the ship arrives full and leaves empty, we select the first of these as the end of the trip.

We check the performance of the algorithm by comparing the duration of some frequent trips, with distances provided by https://sea-distances.org/, and find that we match trip durations well.

2.B Construction of Searching Ships

Here we describe the construction of the vector of searching ships $s_t = [s_{t1}, ..., s_{tI}]$ and matches $m_t = [m_{t1}, \cdots, m_{tI}]$, where s_{it} denotes the number of ships in region i and week t that are available to transport a cargo and m_{it} the realized matches in region i and week t. To construct s_{it} we consider all ships that ended a trip (loaded or ballast) in region i and week $t-1$. We exclude the first week post arrival in the region to account for loading/unloading times (on average (un)loading takes 3-4 days but the variance is large; removing one week will tend to underestimate port wait times). To construct m_{it}, we consider the number of ships that began a loaded trip from region i in week t.

2.C Additional Figures and Tables

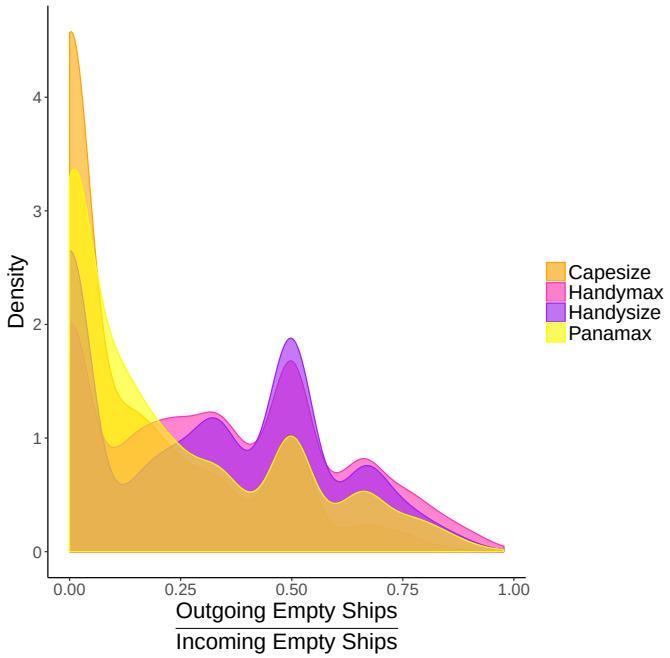

Figure 2.14: Histogram of the ratio of outgoing empty, over incoming empty and loading ships in net exporting countries, by ship type.

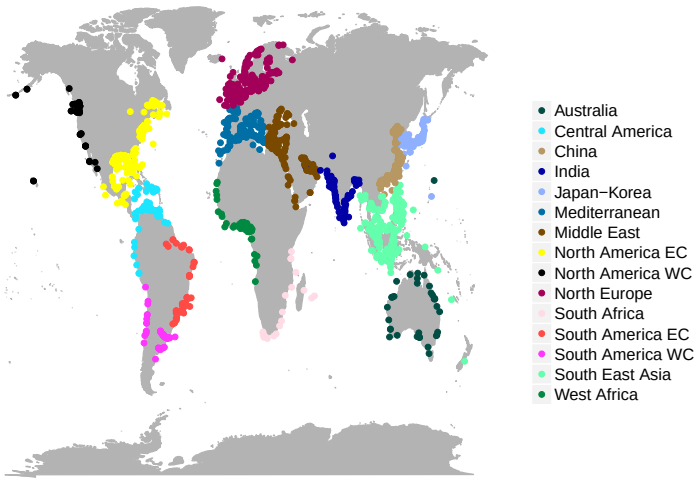

Figure 2.15: Definition of regions. Each color depicts one of the 15 geographical regions.

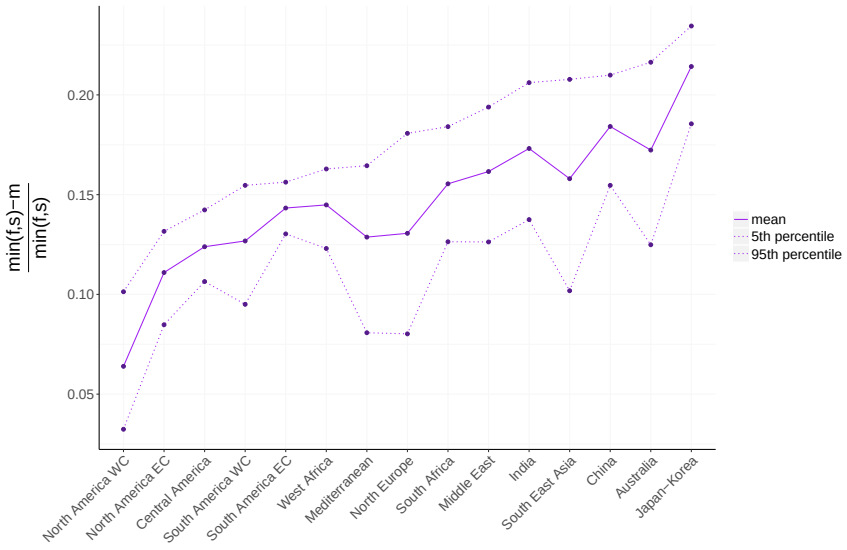

Figure 2.16: Average weekly share of unrealized matches due to search frictions, with confidence intervals from 500 bootstrap samples.

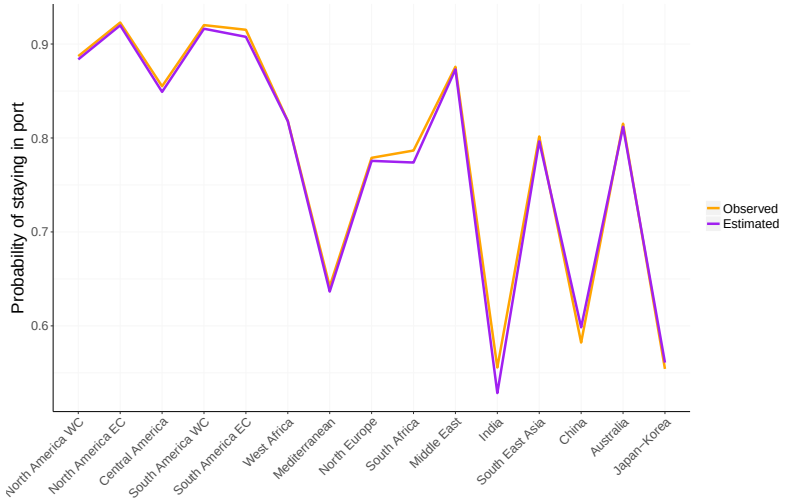

Figure 2.17: Observed and estimated probability of waiting in port P_{ii}.

	N	R^2	Joint Significance	$\frac{s}{m}$
North America West Coast	193	0.146	0	2.706
North America East Coast	200	0.17	0.013	3.172
Central America	199	0.272	0	3.451
South America West Coast	198	0.246	0	2.913
South America East Coast	200	0.269	0	4.083
West Africa	200	0.261	0	5.862
Mediterranean	200	0.358	0	4.244
North Europe	200	0.23	0	3.577
South Africa	200	0.083	0.01	2.862
Middle East	200	0.147	0.001	3.86
India	200	0.12	0.018	8.58
South East Asia	200	0.18	0.005	3.334
China	200	0.177	0	6.194
Australia	187	0.17	0.008	2.457
Japan-Korea	200	0.16	0.003	5.311

Table 2.4: Test for search frictions. Regressions of the number of matches in each region on the unpredictable component of weather conditions in the surrounding seas. For each region we use weeks in which there are at least twice as many ships as matches. The first column reports the number of observations; the second column the R^2; the third column joint significance; and the last column the average ratio between matches and ships in each region. To proxy for the unpredictable component of weather, we divide the sea surrounding each region into 8 different zones (North East, South East, South West and North West both close to the coast and in open sea), and we use the speed of the horizontal (E/W) and vertical (N/S) component of wind in each zone to proxy for weather conditions. Finally, we run a VAR regression of these weather variables on their lag component and season fixed effects and use the residuals, together with their squared term, as independent variables in the regression.

	N	R^2	Joint Significance
North America West Coast	200	0.101	0.004
North America East Coast	200	0.106	0.0002
Central America	200	0.175	0.0007
South America West Coast	200	0.418	0
South America East Coast	200	0.178	0
West Africa	200	0.138	0.0001
Mediterranean	200	0.181	0
North Europe	200	0.138	0.0003
South Africa	200	0.066	0.064
Middle East	200	0.162	0.0012
India	200	0.157	0.0001
South East Asia	200	0.081	0.0008
China	200	0.176	0
Australia	200	0.049	0.02
Japan-Korea	200	0.036	0.12

Table 2.5: First Stage, Matching Function Estimation. Regressions of the number of ships in each region on the unpredictable component of weather conditions in the surrounding seas. The first column reports the number of observations; the second column the R^2; the third column joint significance. To proxy for the unpredictable component of weather, we divide the sea surrounding each region into 8 different zones (North East, South East, South West and North West both close to the coast and in open sea), and we use the speed of the horizontal (E/W) and vertical (N/S) component of wind in each zone to proxy for weather conditions. Finally, we run a VAR regression of these weather variables on their lag component and season fixed effects and use the residuals, together with their squared term, as independent variables in the regression.

2.D Proof of Proposition 2.4.2

We first derive (2.18) and (2.19). Suppose $f_{it}, s_{it}, s^w_{ijt}$ approach f_i, s_i, s^w_{ij} as $t \to \infty$. Then (2.2) becomes:

$$s_i = (s_i - m_i(s_i, f_i)) P_{ii} + \sum_{j \neq i} \xi_{ji} s^w_{ji} \tag{2.28}$$

while for a ship traveling from j to i, (2.3) becomes:

$$s^w_{ji} = (1 - \xi_{ji}) s^w_{ji} + P_{ji} s_j + (G_{ji} - P_{ji}) m_j(s_j, f_j) \tag{2.29}$$

or

$$\xi_{ji} s^w_{ji} = P_{ji} s_j + (G_{ji} - P_{ji}) m_j = P_{ji}(s_j - m_j) + G_{ji} m_j$$

where $m_i = m_i(s_i, f_i)$. Summing this with respect to $j \neq i$ we obtain:

$$\sum_{j \neq i} \xi_{ji} s^w_{ji} = \sum_{j \neq i} P_{ji}(s_j - m_j) + \sum_{j \neq i} G_{ji} m_j$$

and replacing in (2.28) we get (2.18).

Equation (2.19) is a direct consequence of (2.1) and (2.14).

The steady state equations (2.18) and (2.19) have a fixed point over a properly defined subset of $\mathbb{R}^{2I}$, by the Leray-Schauder-Tychonoff theorem (Bertsekas and Tsitsiklis (2015)) which states that if X is a non-empty, convex and compact subset of $\mathbb{R}^{2I}$ and $h : X \to X$ is continuous, then h has a fixed point. Indeed, let $h : \mathbb{R}^{2I} \to \mathbb{R}^{2I}$, $h = (h^s, h^f)$ with:

$$h^s_i(s, f) = \sum_{j=1}^{I} P_{ji}(s, f)(s_j - m_j(s_j, f_j)) + \sum_{j \neq i} G_{ji} m_j(s, f)$$

$$h^f_i(s, f) = \delta(f_i - m_i(s_i, f_i)) + \mathcal{E}_i \sum_{j \neq 0, i} \tilde{G}_{ij}(s, f)$$

for $i = 1, ..., I$. Let $X = \prod_{i=1}^{I} [0, \mathcal{E}_i/(1-\delta)] \times \Delta s$, where $\Delta s = \left\{ s_i \geq 0 \ : \ \sum_{i=1}^{I} s_i \leq S \right\}$. X is nonempty, convex and compact, while h is continuous on X. We assume that the matching function is such that λ, λ^f are zero at the origin and continuous. It remains to show that $F(X) \subseteq X$. Let $(s, f) \in X$. Then, $f_i \leq \mathcal{E}_i/(1-\delta)$ and $\sum_{i=1}^{I} s_i \leq S$. Now,

$$h_i^s(s, f) = \sum_{j=1}^{I} P_{ji}(s, f)(s_j - \lambda_j(s_j, f_j) s_j) + \sum_{j \neq i} G_{ji} \lambda_j(s, f) s_j$$

or

$$h_i^s(s, f) = \sum_{j=1}^{I} s_j [P_{ji}(s, f)(1 - \lambda_j(s_j, f_j)) + G_{ji} \lambda_j(s, f)]$$

where let $G_{ii} = 0$ (no inter-region trips). Summing over i gives:

$$\sum_{i=1}^{I} h_i^s(s, f) = \sum_{j=1}^{I} s_j \left[\sum_{i=1}^{I} P_{ji}(s, f)(1 - \lambda_j(s_j, f_j)) + \sum_{i=1}^{I} G_{ji} \lambda_j(s, f) \right]$$

or

$$\sum_{i=1}^{I} h_i^s(s, f) = \sum_{j=1}^{I} s_j [1 - \lambda_j(s_j, f_j) + \lambda_j(s, f)] \leq S$$

Hence $h_i^s(s, f) \in \Delta s$.

Finally, consider h^f; since $m_i \geq 0$, we have

$$h_i^f \leq \delta f_i + \mathcal{E}_i \sum_{j \neq 0, i} \tilde{G}_{ij}(s, f) \leq \delta f_i + \mathcal{E}_i \leq \delta \frac{\mathcal{E}_i}{1-\delta} + E_i = \frac{\mathcal{E}_i}{1-\delta}$$

Hence $h_i^f(s, f) \in [0, \mathcal{E}_i/(1-\delta)]$.

2.E Proof of Proposition 2.5.1

(i) Following Matzkin (2003), two matching functions $m(\cdot)$ and $\tilde{m}(\cdot)$ are observationally equivalent if there exists a strictly increasing and differentiable function $g(\cdot)$ such

that:

$$\tilde{m}(s, f) = m(s, g(f))$$

Let $\lambda > 0$ and fix $\bar{s}, \bar{f}$. Then

$$\tilde{m}(\lambda\bar{s}, \lambda\bar{f}) = \lambda\tilde{m}(\bar{s}, \bar{f}) = \lambda\bar{m}$$

Furthermore,

$$\tilde{m}(\lambda\bar{s}, \lambda\bar{f}) = m(\lambda\bar{s}, g(\lambda\bar{f})) = \lambda m(\bar{s}, \frac{1}{\lambda}g(\lambda\bar{f}))$$

Therefore,

$$\bar{m} = \tilde{m}(\bar{s}, \bar{f}) = m(\bar{s}, \frac{1}{\lambda}g(\lambda\bar{f}))$$

Invertibility implies that $\bar{f} = \tilde{m}^{-1}(\bar{s}, \bar{m})$ and $\frac{1}{\lambda}g(\lambda\bar{f}) = m^{-1}(\bar{s}, \bar{m})$, or

$$g(\lambda\bar{f}) = \lambda m^{-1}(\bar{s}, \bar{m})$$

Differentiate with respect to λ to obtain $\bar{f}g'(\lambda\bar{f}) = m^{-1}(\bar{s}, \bar{m})$, which for $\lambda = 1$ becomes $g'(\bar{f})\bar{f} = m^{-1}(\bar{s}, \bar{m}) = g(\bar{f})$. Therefore, the Euler condition is satisfied and $g(\cdot)$ is homogeneous of degree 1. Since $g(\cdot)$ is a function of a real variable, the only possibility is $g(f) = cf$ with $c > 0$, a constant. Finally, we use the a priori knowledge of the point $m^* = m(s^*, f^*)$ to establish that $c = 1$. Indeed, by definition, $m(s^*, f^*) = \tilde{m}(s^*, f^*) = m^*$. But also, $m(s^*, cf^*) = \tilde{m}(s^*, f^*)$. Therefore, $cf^* = f^*$ and since $f^* \neq 0$, $c = 1$.

(ii) Conditional on η, s is a function of z which in turn is by assumption independent from f. It follows that s and f are conditionally independent given η. At a point ϕ we have that:

$$F_{f|\eta}(\phi) = \Pr(f \leq \phi|\eta) = \Pr(f \leq \phi|\eta, s) = \Pr(m(s, f) \leq m(s, \phi)|\eta, s) = F_{m|s,\eta}(m(s, \phi))$$

Hence,

$$m(s,\phi) = F^{-1}_{m|s,\eta}(F_{f|\eta}(\phi))$$

which we integrate over η to obtain the result. Let $\phi = \frac{\phi}{f^*} f^*$. Then,

$$F_{f|\eta}(\phi) = F_{m|s=\frac{\phi}{f^*}s^*,\eta}\left(m\left(\frac{\phi}{f^*}s^*, \frac{\phi}{f^*}f^*\right)\right) = F_{m|s=\frac{\phi}{f^*}s^*,\eta}\left(\frac{\phi}{f^*}m^*\right)$$

2.F Proof of Lemma 2.5.2

Fix θ. Let $\phi_{ij} = \frac{1}{1-\beta(1-\xi_{ij})}$. The map $T_\theta(U)$ is differentiable with respect to U with Jacobian:

$$\frac{\partial T_\theta(U)}{\partial U} = \beta\left(DG + (I-D)\,P\right) \odot Z \tag{2.30}$$

where D is a diagonal matrix with λ_i it's i diagonal entry; P is the matrix of choice probabilities, G is the matrix of matched trips, Z is an $L \times L$ matrix whose (i,j) element is $\phi_{ij}\xi_{ij}$ and $\odot$ denotes the pointwise product. Indeed, the (i,j) entry of $\frac{\partial T}{\partial U}$ is

$$\left(\frac{\partial T}{\partial U}\right)_{ij} = 1\{i=j\} - \beta\lambda_i G_{ij}\xi_{ij}\phi_{ij} - (1-\lambda_i)\frac{\partial J_i}{\partial U_j}$$

Now,

$$\frac{\partial J_i}{\partial U_j} = \frac{1}{e^{\frac{\beta U i}{\sigma}} + \sum_k e^{\frac{W_{ik}}{\sigma}}} e^{\frac{W_{ij}}{\sigma}} \frac{\partial W_{ij}}{\partial U_j} = \beta P_{ij}\xi_{ij}\phi_{ij}$$

and thus

$$\left(\frac{\partial T}{\partial U}\right)_{ij} = 1\{i=j\} - \beta\left(\lambda_i G_{ij} + (1-\lambda_i)\,P_{ij}\right)\xi_{ij}\phi_{ij}$$

which in matrix form becomes (2.30) (as a convention set $\xi_{ii} = 1$). Let $H = (DG + (I-D)P) \odot Z$. Take $||H|| = \max_i \sum_j |H_{ij}|$. Note that G, P are stochastic matrices and the diagonal matrix D is positive with entries smaller than 1. Thus $DG+(I-D)P$

is stochastic. It is also true that $0 < \xi_{ij}\phi_{ij} \leq 1$. Thus,

$$\sum_j |H_{ij}| = \sum_j \left(\lambda_i G_{ij} + (1-\lambda_i) P_{ij}\right) \xi_{ij}\phi_{ij} \leq \sum_j \left(\lambda_i G_{ij} + (1-\lambda_i) P_{ij}\right) \leq 1$$

and therefore $||H|| \leq 1$. We deduce that $||\frac{\partial T_\theta(U)}{\partial U}|| \leqq \beta < 1$. The mean value theorem then implies

$$||T_\theta(U) - T_\theta(U')|| \leq \beta ||U - U'||$$

2.G Identification of Ship Port and Sailing Costs

Given the choice probabilities $P_{ij}(\theta)$, the parameters $\{\frac{c_{ij}^s}{\sigma}, \frac{c_i^u}{\sigma}, \frac{1}{\sigma}\}$ satisfy a $(I^2 - I) \times (I^2 + 1)$ linear system of equations of full rank $I^2 - I$. Hence, $(I + 1)$ additional restrictions are required for identification. Let $\phi_{ij} = \frac{1}{1-\beta(1-\xi_{ij})}$. The Hotz and Miller (1993) inversion states:

$$\sigma \log \frac{P_{ij}}{P_{ii}} = W_{ij}(\theta) - \beta U_i(\theta)$$

Substituting from (2.4)-(2.5) we obtain:

$$\sigma \log \frac{P_{ij}}{P_{ii}} = -\phi_{ij} c_{ij}^s + \beta \xi_{ij} \phi_{ij} U_j(\theta) - \beta U_i(\theta) \tag{2.31}$$

It also holds that (see Kalouptsidi et al. (2016))):

$$\log P_{ij} = \frac{W_{ij}}{\sigma} - \frac{J_i}{\sigma} + \gamma^{euler}$$

or:

$$\sigma \log P_{ij} = -\phi_{ij} c_{ij}^s + \beta \xi_{ij} \phi_{ij} U_j(\theta) - J_i + \sigma\gamma^{euler} \tag{2.32}$$

and

$$\sigma \log P_{ii} = \beta U_i(\theta) - J_i + \sigma\gamma^{euler} \tag{2.33}$$

Now, replace W_{ij} from (2.32) into the definition of U, (2.5) to get:

$$U_i(\theta) = -c_i^u + \lambda_i \tau_i + \sigma \lambda_i \sum_{j \neq i} G_{ij} \log P_{ij} - \sigma \lambda_i \gamma^{euler} + J_i$$

where $\tau_i \equiv E_{v,j} \tau_{ijv} = \sum_{j \neq i} G_{ij} \tau_{ij}$. Substitute J_i from (2.33):

$$U_i(\theta) = -\frac{1}{1-\beta} c_i^u + \frac{\sigma}{1-\beta} \left((1-\lambda_i)\gamma^{euler} + \lambda_i \sum_{j \neq i} G_{ij} \log P_{ij} - \log P_{ii} \right) + \frac{1}{1-\beta} \lambda_i \tau_i$$

so that given the CCP's, U_i is an affine function of c^u and σ. Next, we replace this into the Hotz and Miller (1993) inversion (2.31) to obtain:

$$c_{ij}^s = \frac{\beta}{\phi_{ij}(1-\beta)} c_i^u - \frac{\beta}{1-\beta} \xi_{ij} c_j^u +$$

$$+\sigma \left(\frac{\beta}{1-\beta} \xi_{ij} \left[(1-\lambda_j)\gamma^{euler} + \lambda_j \sum_{l \neq j} G_{jl} \log P_{jl} - \log P_{jj} \right] \right) -$$

$$-\sigma \left(\frac{\beta}{1-\beta} \frac{1}{\phi_{ij}} \left[(1-\lambda_i)\gamma^{euler} + \lambda_i \sum_{l \neq i} G_{il} \log P_{il} - \log P_{ii} \right] \right) -$$

$$-\frac{\sigma}{\phi_{ij}} \log \frac{P_{ij}}{P_{ii}} + \frac{\beta}{1-\beta} \xi_{ij} \lambda_j \tau_j - \frac{\beta}{(1-\beta)\phi_{ij}} \lambda_i \tau_i$$

Note that

$$\frac{1}{\phi_{ij}(1-\beta)} = \frac{1-\beta(1-\xi_{ij})}{1-\beta} = 1 + \frac{\beta \xi_{ij}}{1-\beta}$$

and set $\rho_{ij} = \frac{\beta \xi_{ij}}{1-\beta}$, then $\frac{1}{(1-\beta)\phi_{ij}} = 1 + \rho_{ij}$.

We divide by σ:

$$\frac{c_{ij}^s}{\sigma} = (1+\rho_{ij}) \frac{c_i^u}{\sigma} - \rho_{ij} \frac{c_j^u}{\sigma} - \left[\beta (1+\rho_{ij}) \lambda_i \tau_i - \rho_{ij} \lambda_j \tau_j \right] \frac{1}{\sigma} +$$

$$+\rho_{ij} \left[(1-\lambda_j) \gamma^{euler} + \lambda_j \sum_{l \neq j} G_{jl} \log P_{jl} - \log P_{jj} \right]$$

$$-\beta(1+\rho_{ij})\left[(1-\lambda_i)\,\gamma^{euler}+\lambda_i\sum_{i\neq j}G_{il}\log P_{il}-\log P_{ii}\right]$$

$$-\frac{1}{\phi_{ij}}\log\frac{P_{ij}}{P_{ii}}$$

This is a linear system of full rank in the parameters $\{\frac{c^s_{ij}}{\sigma},\frac{c^u_i}{\sigma},\frac{1}{\sigma}\}$, since $\frac{c^s_{ij}}{\sigma}$ can be expressed with respect to $\{\frac{c^u_i}{\sigma},\frac{1}{\sigma}\}$.

2.H Algorithm for computing the steady state equilibrium

Here, we describe the algorithm employed to compute the steady state of our model to obtain the counterfactuals of Section 2.7.

1. Make an initial guess for $\{s^0, f^0, U^0\}$.

2. At each iteration m, inherit $\{s^m, f^m, U^m\}$

   (a) Update the ship's and exporter's optimal policies by repeating the following steps K times.[50]

      i. Solve for W^{m+1} from:

$$W^{m+1}_{ij}=\frac{-c^s_{ij}+\xi_{ij}\beta U^m_j}{1-\beta\,(1-\xi_{ij})}$$

      ii. Update J^{m+1} from:

$$J^{m+1}_i=\sigma\log\left(\exp\frac{\beta U^m_i}{\sigma}+\sum_{j\neq i}\exp\frac{W^m_{ij}}{\sigma}\right)+\sigma\gamma^{euler}$$

[50] K is chosen to accelerate the convergence in the spirit of standard modified policy iteration methods.

iii. Compute the equilibrium prices using

$$\tau_{ijv}^{m} = \frac{\gamma\left(1-\beta\delta\left(1-\lambda_i^{f,m}\right)\right)}{1-\beta\delta\left(1-\gamma\lambda_i^{f,m}\right)}\left(J_i^{m+1}-W_{ij}^{m+1}\right)+\frac{(1-\gamma)(1-\beta\delta)}{1-\beta\delta\left(1-\gamma\lambda_i^{f}\right)}\mu_{ij}$$

iv. Update $\tilde{G}$:

$$\tilde{G}_{ij}^{m+1} \equiv \frac{\exp\left(\frac{\lambda_i^f\left(\mu_{ij}-\tau_{ij}^m\right)}{1-\beta\delta\left(1-\lambda_i^f\right)}-\kappa_{ij}\right)}{1+\sum_{l\neq i}\exp\left(\frac{\lambda_i^f\left(\mu_{ij}-\tau_{ij}^m\right)}{1-\beta\delta\left(1-\lambda_i^f\right)}-\kappa_{il}\right)}$$

v. Update U^m:

$$U_i^{m+1} = -c_i^u + \lambda_i E_{v,j}\tau_{ijv} + \lambda_i\sum_{j\neq i}\left(\frac{\tilde{G}_{ij}^{m+1}}{1-\tilde{G}_{i0}^{m+1}}\right)W_{ij}^{m+1}+(1-\lambda_i)J_i^{m+1}$$

vi. Obtain the ships ballast choices $\left(P_{ij}^{m+1}\right)_{i=1:I,j=1:I}$

3. Update to $\{\tilde{s}^{m+1},\tilde{f}^{m+1}\}$ from:

$$\tilde{f}_i^{m+1} = \delta_i\left(f_i^m - m_i^m\right)+\mathcal{E}_i\left(1-\tilde{G}_{i0}^{m+1}\right)$$

and

$$\tilde{s}_i^{m+1} = \sum_j P_{ji}^{m+1}\left(s_j^m-m_j^m\right)+\sum_j\frac{\tilde{G}_{ij}^{m+1}}{1-\tilde{G}_{i0}^{m+1}}G_{ji}^{m+1}m_j^m$$

4. If $\|\tilde{s}^{m+1}-s^m\|<\epsilon$, $\left\|\tilde{f}^{m+1}-f^m\right\|<\epsilon$ and $\|U^{m+1}-U^m\|<\epsilon$, stop, otherwise update freights and ships as follows:

$$\begin{aligned} s^{m+1} &= \alpha s^m + (1-\alpha)\,\tilde{s}^{m+1} \\ f^{m+1} &= \alpha f^m + (1-\alpha)\,\tilde{f}^{m+1}, \end{aligned}$$

where α is a smoothing parameter.

Chapter 3

The Impact of Oil Prices on World Trade[1]

3.1 Introduction

How significant an effect do oil prices have on trade flows? Many have claimed that oil price spikes have the potential to put a break on world trade by increasing transport costs.[2] Indeed, oil prices determine ship fuel costs, which constitute the core (variable) cost component of the transportation sector. In this paper, we compute the world trade elasticity with respect to oil prices. We are able to isolate the impact of oil shocks on transportation costs (vs. other channels such as production input costs or income effects) through the use of a simple structural trade model that explicitly incorporates the transportation sector. We show that modeling the transport sector is

[1]The paper on which this chapter is based is co-authored with Myrto Kalouptsidi and Theodore Papageorgiou.

[2]For instance, a report by The Economist (Wood, 2008) states "For the countries of Asia, where the price of transporting goods to the West has a significant impact on their attractiveness as a manufacturing location, these [the rising oil prices] are serious issues. In an era of high and volatile oil prices, is the world not as flat as has been suggested?". Similarly, in 2008 Paul Krugman writes in his blog "higher fuel prices are putting the brakes on globalization: if it costs more to ship stuff, there will be less shipping". In his book "Why your world is about to get a whole lot smaller: oil and the end of globalization", Rubin (2009) argues that spiking oil prices will curtail long-distance shipping and traveling.

crucial in understanding the mechanism behind the oil price pass-through to exporters and the patterns of the estimated elasticity.

We focus on oceanic shipping, which accounts for the large majority of international trade, and in particular on trade in bulk commodities, such as minerals, grain, ores and chemicals, which in turn accounts for about half of all seaborne trade in tonnage (UNCTAD, 2015); however, our findings may hold more generally. Bulk ships are often thought of as the "taxis of the ocean", as the industry's structure and operation resembles that of taxicabs. For this sector of the transportation industry, we are able to collect data on shipping contracts between shipowners and exporters that correspond to specific trips; each transaction includes the shipping price as well as the origin and destination of the trip. We also obtain AIS ship movement data that inform us on ships' sequences of loaded and empty trips. This dataset was first used in our prior work, Brancaccio et al. (2017), henceforth BKP.

We employ the dynamic spatial search model built in BKP. This model is in the spirit of Mortensen and Pissarides (1994) and Lagos (2000) and centers on the behavior of ships and exporters. Ships are homogeneous and contract with exporters for individual trips. Fuel costs are the main variable cost of ships and they are captured in our setup through ships' cost of sailing. When a ship and an exporter meet in a particular region of the world, they bargain over the price to transport the exporter's cargo from this location to the exporter's desired destination. The ship then travels the distance and restarts in the destination by searching for a new cargo. In every period, if a ship does not find a match with an exporter, it makes a choice: it can either wait at its current location, or it can choose a destination to ballast to (i.e. travel empty). Exporters have one cargo to ship. Potential exporters decide whether to export or not, as well as which destination to export to. Once this entry decision is made, they wait at port until they can match with a ship that will transport their cargo. The model parameters (the matching function describing the meeting process

between exporters and ships, the valuations and entry costs of exporters and the ship port costs) are estimated using our datasets as in BKP.

A key feature of the data that our model is designed to capture, is the impact of trade imbalances on shipping prices and trade flows. Indeed, as discussed in BKP, most countries are either net importers or net exporters of bulk commodities. This asymmetry is also reflected in shipping prices: it is more expensive to hire a ship when the destination is a net importing country, as the ship's profitability is low there (it has to either wait for a cargo or incur the fuel costs of ballasting elsewhere). These features are important in determining the pass-through of oil prices to shipping prices, as well as the world trade elasticity with respect to oil prices.

We compute the trade elasticity by considering how changes in the fuel costs pass-through to shipping prices and in turn affect trade flows. Our estimated trade elasticity is 0.35 at the average observed fuel cost level. However, it depends crucially on the level of the oil price and in fact ranges between 0.1 and 1.2. We also estimate that the pass-through of fuel costs to exporters is low, as the elasticity of prices with respect to fuel costs equals 0.17 at the average observed fuel cost. This suggests that ship variable costs are not a good approximation of transport costs and that modeling the transport sector is important in understanding trade flows and trade costs.

A striking feature of the estimated trade elasticity is its pronounced asymmetry in low and high levels of fuel costs. Indeed, the elasticity gets steeper as the fuel cost increases, while it plateaus as the fuel cost decreases. This asymmetry is generated by the equilibrium of the transportation sector and in particular, the changes in the relative bargaining positions of ships and exporters. Naturally, as the fuel cost declines, trade increases, since fuel costs are a key input in transportation costs. In addition, however, fuel costs are a key determinant of ships' relative bargaining position. As the fuel cost declines, the world becomes more "flat", since distance matters less, and ships can reallocate cheaply across different regions. Therefore,

they are less "tied" to their current location and are able to extract higher prices than otherwise. This dampens the price decline and mutes the increase in trade disproportionally at low fuel costs. In contrast, when fuel costs are high, it is costlier for ships to change locations and exporters at their current location have a stronger bargaining position, leading to large increases in trade. This effect is particularly pronounced in net exporting regions, where the high likelihood of finding a load makes ships almost certain to stay put when oil prices are high. As a result, as the world becomes flat under low fuel costs, the increase in trade shrinks because of the ships' strong bargaining position.

To further document this mechanism, we explore a relevant testable prediction of our model. When fuel costs decline and distance is of less importance, ship values tend to equalize over space. A ship in an unloading region is now not much worse off than a ship in a loading region, as ballasting from the former to the latter is cheaper. As the dispersion in ship value functions declines, so does the dispersion in shipping prices which depend on the ship's value at the destination. We test this using the observed shipping prices and oil prices and find that indeed as oil prices increase, shipping prices equalize across space.

Finally, we use our estimates to assess how much recent trends in fuel efficiency of ship design have affected trade flows. Ship design is affected by a number of factors over time, such as long-term trends in the shipbuilding industry, technological improvements and environmental policies. For instance, the 1980s witnessed large improvements in fuel efficiency (following the oil price shocks of the 1970s), while the last 10 to 20 years witnessed a reversion, with fuel efficiency deteriorating by about 13%. We compute that the efficiency gains achieved in the 1980s by world shipyards led to a decline in shipping prices of 5.5% and an increase in trade by 12%. On the other hand, the recent deterioration in fuel efficiency design resulted in a 5.6% reduction in world trade. These calculations showcase the policy relevance of our

estimated trade elasticity with respect to fuel costs, as a number of environmental regulations imposing fuel efficiency targets are currently discussed by international organizations.

The literature quantifying the impact of oil prices on trade, as well as transport costs, is relatively thin. von Below and Vezina (2016) (see also work cited therein) employ a gravity equation that incorporates oil prices to measure their impact on trade through transportation costs. The estimated elasticity is fairly large, between -1.2 and -1.8. A challenge with this approach is that it is potentially difficult to separate the impact of oil shocks operating through transport costs, as opposed to other facets of the economy (the authors do adopt an IV approach to handle this). Hummels (2007) measures the elasticity of freight costs with respect to fuel costs and estimates it to about 0.2 to 0.3, while a few recent papers have explored the elasticity of trade with respect to freight rates, as well as the role of the transportation sector overall (e.g. Hummels (2007); Asturias (2016); Wong (2017), BKP). For instance, Wong (2017) estimates the trade elasticity with respect to shipping prices to about -3 and Limao and Venables (2001) report a similar estimate.[3] Bridgman (2008) considers a trade model with an energy-using transportation sector to investigate the 1970s and 1980s oil shocks and finds that they had a significant impact on world trade. Rubin (2009) argues that a shortage in oil and increasing oil prices will limit globalization.

The approach adopted here to measure the trade elasticity is different from most of the existing literature; we employ our structural model to isolate the impact of oil shocks only through the transportation sector rather than other channels (e.g. higher input costs, income effects, or correlated changes in demand). In addition, our model guides us in understanding the shape of the elasticity, as well as the mechanism of the oil price pass-through to exporters.

[3]North (1958) and Estevadeordal et al. (2003) also find that changes in transportation prices have been historically important determinants of world trade.

Finally, there is a strand of literature that investigates how oil prices affect trade more broadly, besides its impact through transport prices (e.g. Backus and Crucini (2000), Kilian et al. (2009), Chen and Hsu (2012) and others). In addition, a number of papers in the field of maritime economics have explored ships' optimal speed choices, as well as fuel efficiency, emissions and environmental policies on energy efficiency in different ship segments (e.g. Adland and Jia (2016a), Adland and Jia (2016b), Adland et al. (2017) and Adland et al. (2017)). Lastly, more broadly, our paper contributes to the large literature on trade costs (e.g. Anderson and Van Wincoop, 2003) and geography (e.g. Krugman, 1991), while our methodology borrows from the theoretical and empirical literature on dynamic models (Rust (1987); Hopenhayn (1992); Ericson and Pakes (1995)).

The paper proceeds as follows. In Section 3.2 we discuss the key features of the oceanic bulk shipping industry and the data. In Section 3.3 we present the model, as well as the procedure employed to estimate it. In Section 3.4 we present our main results on the trade elasticity with respect to fuel costs and the oil price pass-through to exporters, while in Section 3.5 we investigate our setup's prediction regarding the relationship between price dispersion and oil prices. In Section 3.6 we ask how much the recent improvements in ship fuel efficiency have contributed to trade growth. Finally, in Section 3.7 we conclude.

3.2 Industry and Data

Bulk shipping involves large oceanic carrier vessels (larger than 10,000 DWT capacity) that carry mostly commodities and raw materials, such as grain, iron ore, steel, coal, chemicals, etc. The industry is very unconcentrated with a large number of small firms. These ships operate much like taxi cabs: a shipowner contracts with a cargo owner for a specific trip; the ship is filled up with this exporter's load and

it delivers the cargo at the agreed upon destination. The ship then restarts in that destination by looking for a new contract. Similar to taxi cabs, bulk shipping services are considered fairly homogeneous. Unlike taxis however, prices are not regulated; they are negotiated between the shipowner and the exporter, and mediated by one or multiple shipbrokers.

We exploit two main databases. The first consists of a sample of contracts between shipowners and exporters obtained from Clarksons Research. Each observation is a contract for a trip and it specifies the origin and destination of the trip, the loading and signing dates, the ship, the shipowner and cargo owner and finally the price. We use this dataset to obtain information on trip prices.

The second dataset, obtained from ExactEarth Ltd, contains ship movements collected from satellites (we observe the location of about half of the world fleet every six minutes). The AIS data also report the ship's draft (i.e. the distance between the bottom of the ship's hull and the waterline) which allows us to distinguish loaded from empty movements. We employ this dataset to construct ship movement histories of loaded and empty trips. For a more detailed description of the industry we refer the reader to BKP and Kalouptsidi (2014); BKP also provides more details on the data, as well as summary statistics and notable data patterns.

A prevailing feature of the data is the large trade imbalances and their impact on shipping prices. Indeed, as documented in BKP, most countries are either large net importers or large net exporters of the commodities carried in bulk vessels. China and India are the biggest importers, while Brazil, Australia and North America are the biggest exporters. These trade imbalances translate into asymmetric ship hiring rates at different regions of the world: although a ship in Brazil is very likely to find a cargo, a ship in China is much less so. As a result, a ship would much rather unload a cargo in Brazil than China, as her options are much better in Brazil and this is reflected in the prices the ship agrees upon. Indeed, the price to unload in

China is substantially higher than the price to unload in Brazil. Shipping prices exhibit pronounced asymmetries, reflecting the world's natural geography (distances and natural inheritance) and its impact on ship profitability at different regions of the world.

3.3 Model

We next provide a description of the dynamic spatial search model of the global shipping industry introduced in BKP. We refer the interested reader to BKP for further details.

3.3.1 Environment

Time is discrete. There are I regions in the world and two types of agents: exporters (or freights) and ships. Both exporters and ships are risk neutral and share a discount factor β. Let f_i denote the freights awaiting transportation in region i. Freights are heterogeneous in (i) their destination, j (endogenized below); (ii) their valuation, v, denoting the value of the cargo upon delivery to the destination.

Ships are homogeneous and carry at most one freight at a time. Every period a ship is either sailing towards destination j, full or empty, at a per-period sailing cost c^s_{ij}, which is primarily the fuel cost; or it is waiting in port i at cost c^u_i. A ship sailing from i to j arrives at its destination with probability ξ_{ij}, so that the average trip duration equals $1/\xi_{ij}$.

Ships at port i meet exporters originating from port i randomly. Every period the number of matches in location i is determined by a matching function $m_i(f_i, s_i)$, where s_i is the number of unmatched ships in location i, and f_i is the number of unmatched freights. Let $\lambda_i = m_i/s_i$ denote the probability that a ship meets a freight and $\lambda^f_i = m_i/f_i$ the probability that a freight meets a ship. We follow the bulk of the

search literature, and assume that the surplus of a meeting is split via generalized Nash bargaining. This determines the price τ_{ijv} that a freight with valuation v pays the ship in order to be transported from i to j. Let $\gamma \in (0,1)$ denote the freight's bargaining power.[4]

Timing is as follows:

1. Ships and freights match.

2. Unmatched ships draw additive iid preference shocks $\epsilon = [\epsilon_1, ..., \epsilon_I] \in \mathbb{R}^I$ from a Type I extreme value distribution and decide whether to (i) stay in their current region and wait for freight; or (ii) ballast toward some destination j. Each preference shock is associated with one of these choices.

3. Ships already traveling from i to j, arrive with probability ξ_{ij}. Unmatched ships that decided to ballast begin traveling to their chosen destination. Existing unmatched exporters disappear with probability $1 - \delta$.

4. In each region i, $\mathcal{E}_i$ potential exporters decide whether and to which destination to export to. The exporters that do enter the market, pay an entry cost κ_{ij}, draw their valuations v from a distribution F_{ij}^v, and join the pool of unmatched exporters the following period.

3.3.2 Equilibrium

In this section, we derive firm behavior. We focus on the steady state equilibrium of this framework.

[4]We implicitly assume that freight valuations, v, are large enough so that all meetings are converted into matches, i.e. the match surplus is always positive. Given that in this context, valuations are an order of magnitude greater than the shipping price, this assumption seems quite reasonable.

Ships The value of a ship traveling from i to j, W_{ij}, is given by,

$$W_{ij} = -c_{ij}^s + \xi_{ij}\beta U_j + (1 - \xi_{ij})\,\beta W_{ij} \tag{3.1}$$

i.e. the ship pays the per period travel cost, c_{ij}^s, then with probability ξ_{ij} it arrives at its destination and begins the following period unmatched at j, and with the complement probability the following period it is still traveling. U_j is the value of a ship that is unmatched in port j at the start of the period and is given by,

$$U_i = -c_i^u + \lambda_i E_{j,v} V_{ijv} + (1 - \lambda_i)\,J_i \tag{3.2}$$

i.e. the ship pays the port cost, c_i^u, and with probability λ_i it meets some freight with destination j and value v, while with the complement probability it remains unmatched and needs to make a decision. We examine each case in turn.

The value of a ship that is matched with a freight with value v and destination j is

$$V_{ijv} = \tau_{ijv} + W_{ij}$$

i.e. the ship receives the shipping price, τ_{ijv}, which is the outcome of Nash bargaining, and immediately begins its journey towards j, obtaining value W_{ij}, defined in (3.1) above.

The value of a ship that remains unmatched in i at the end of the period is

$$J_i(\epsilon) = \max\left\{\beta U_i + \sigma\epsilon_{ii}, \max_{j \neq i} W_{ij} + \sigma\epsilon_{ij}\right\}$$

i.e. the ship can either remain in i and obtain value U_i the following period, defined in (3.2) or it can choose a destination j and ballast there, obtaining value W_{ij}. The terms ϵ_{ij} capture unobserved idiosyncratic ship costs or measurement error and have

standard deviation σ. Finally let $J_i \equiv E_\epsilon J_i(\epsilon)$ denote the "ex ante" value of an unmatched ship, i.e. *before* drawing ϵ_{ij}.

Freights The value of an unmatched freight in market i with destination j and valuation v at the end of the period is

$$J^f_{ijv} = \delta\beta\left(\left(1-\lambda^f_i\right)J^f_{ijv} + \lambda^f_i V^f_{ijv}\right) \tag{3.3}$$

i.e. conditional on surviving (which occurs with probability δ), with probability $1-\lambda^f_i$ it remains unmatched the following period, and with the complement probability it is matched with a ship and receives

$$V^f_{ijv} = v - \tau_{ijv} \tag{3.4}$$

i.e. its value v minus the shipping price, τ_{ijv}.

Shipping Price Nash bargaining implies the surplus sharing condition:

$$\gamma\left(V_{ijv} - J_i\right) = (1-\gamma)\left(V^f_{ijv} - J^f_{ijv}\right)$$

Using this condition and substituting in the value functions, allows us to solve out for the shipping price:

$$\tau_{ijv} = \frac{\gamma\left(1-\beta\delta\left(1-\lambda^f_i\right)\right)}{1-\beta\delta\left(1-\gamma\lambda^f_i\right)}\left(J_i - W_{ij}\right) + \frac{(1-\gamma)\left(1-\beta\delta\right)}{1-\beta\delta\left(1-\gamma\lambda^f_i\right)}v \tag{3.5}$$

The equilibrium price depends positively on v, so that more valuable freights pay higher prices. Moreover, note that the price rises with travel time, $1/\xi_{ij}$, (through the traveler's value W_{ij} defined in equation (3.1)), as well as the the per-period travel costs, c^s_{ij}. It depends negatively on the value of an unmatched ship at the

destination, U_j: freights that are headed towards destinations that offer a high value to ships pay less. From the value of U_i, such destinations may have low port costs, c_i^u, attractive ballast opportunities, high matching probability, λ_i, freights with high values, v and/or freights that are headed to attractive destinations (see equation (3.2)). Finally, the equilibrium price depends on the ship's outside option, i.e. its value if unmatched in i, J_i. Ships matched in regions with high matching probabilities, high value freights or attractive ballast opportunities have higher outside options and as a result command higher prices.

Freight Entry Finally, we characterize the freight entry decision. Every period there are $\mathcal{E}_i$ potential exporters and the value of each is

$$J_i^{ef} = \max\left\{\epsilon_0^f, \max_{j\neq i}\left\{E_v J_{ijv}^f - \kappa_{ij} + \epsilon_j^k\right\}\right\} \tag{3.6}$$

i.e. each one chooses between its outside option of not exporting, in which case the payoff has been normalized to zero, and exporting to one of the $j \neq i$ destinations. If it chooses to export, it pays a destination-specific production and exporting cost, κ_{ij}, it then draws value, v from F_{ij}^v, and obtains the value of an unmatched freight, J_{ijv}^f, given by (3.3). In addition, each potential entrant draws a vector of additive iid choice-specific preference shocks, ϵ^f, from a type I extreme value distribution.

3.3.3 Model Estimation and Results

We now provide a brief description of the estimation strategy followed in BKP. We estimate (i) the matching function and the freights searching for ships, (ii) the port costs, c_i^u, (iii) the distributions of freight values, F_{ij}^v and (iv) the production and exporting costs, κ_{ij}.

Matching Function Estimation In contrast to the majority of empirical work on estimating matching functions, we do not observe all necessary ingredients, $\{s, f, m\}$; indeed, we lack data on one side of the market, namely searching freights.[5] Our goal is thus to use data on the number of ships searching and the number of matches to recover the matching function and the number searching freights.

In order to overcome this lack of data, as well as to avoid imposing functional form assumptions on the matching function, we use results from the literature on non-parametric identification, namely Matzkin (2003). We provide an intuitive description of the approach and refer the reader to BKP for the formal treatment. As is common in the literature, assume that the matching function is strictly increasing in the number of freights. Then, from

$$m_{it} = m_i\left(s_{it}, f_{it}\right)$$

if we knew $m_i\left(.\right)$, as well as m_{it} and s_{it}, we could invert the matching function and recover f_{it}. Conversely, if we observed f_{it}, but not $m_i\left(.\right)$, we could use a nonparametric estimator to recover the matching function. In our case, however, we observe neither $m_i\left(.\right)$, nor f_{it}.

First assume that s_{it} and f_{it} are independent. Then, by leveraging the correlation between m_{it} and s_{it} in the data, we can identify the elasticity of the matching function with respect to s_{it}: since s_{it} and f_{it} are independent, for a given shift in s_{it}, the resulting change in m_{it} is informative about the underlying elasticity. In practice, we relax the independence assumption by using an instrument: we leverage unexpected shocks to sea weather that exogenously shocks the arrivals of ships in a market.

Second assume the matching function exhibits constant returns to scale. This assumption is consistent with the literature on matching function estimation which has

[5]See for instance Petrongolo and Pissarides (2001); two exceptions to the literature are Buchholz (2016) and Frechette et al. (2016).

largely failed to reject constant returns to scale (see e.g. Petrongolo and Pissarides, 2001). This assumption implies that knowing the elasticity of the matching function with respect to s_{it} is sufficient to back out its elasticity with respect to f_{it}. Now that we have essentially recovered the shape of the matching function, we can invert the function to back out the unobserved number of freights, f_{it}, at any point in our sample.

Ship Costs Consider next the ship conditional choice probabilities; the probability that a ship chooses to remain at market i rather than ballast elsewhere is

$$p_{ii} = \frac{\exp(\beta U_i/\sigma)}{\exp(\beta U_i/\sigma) + \sum_{j\neq i} \exp(W_{ij}/\sigma)},$$

while the probability that it chooses to ballast to market j is

$$p_{ij} = \frac{\exp(W_{ij}/\sigma)}{\exp(\beta U_i/\sigma) + \sum_{j\neq i} \exp(W_{ij}/\sigma)}.$$

These model implied probabilities depend on the ship values, U_i and W_{ij}, which in turn depend on travel costs, c^s_{ij} and port costs, c^u_i, as well as prices, τ_{ijv} (see (3.1) and (3.2)).

The travel costs, c^s_{ij} consist primarily of the cost of fuel (Stopford, 2009). We thus calibrate their level using the average weekly price of fuel in our sample, and calculate that weekly travel costs are equal to about \$69,100.[6] We allow travel costs to differ for each pair (i, j) as follows: c^s_{ij} takes one of seven values based on the continent and coast of the origin; we set one of the c^s_{ij} (East Coast of North and South America) to equal to the average fuel cost of 69,100\$ and estimate the rest as described below.[7]

[6]Over our sample period ships face an average fuel price of 470\$ per meter ton. Assuming ships travel at an average speed of 13mph, the average daily consumption of fuel is around 20 tons per day (Stopford (2009)). This adds up to a total expenditure for fuel of \$69,100 per week.

[7]We cluster ports into 15 regions (see Table 3.1). Moreover, the seven groups for c^s_{ij} are: (i) Central America, West Coast Americas; (ii) East Coast Americas; (iii) West and South Africa; (iv) Mediterranean, Middle East and North Europe; (v) India; (vi) Australia and Southeast Asia; (vii)

Note also that since the fuel cost is paid by the exporter when the ship is loaded, we add it to the observed prices; this generates an asymmetry in the cost of traveling full and traveling empty.

We back out the parameters of interest $\left\{c_i^u, c_{ij}^s, \sigma\right\}$ via Maximum Likelihood using ships' observed choice probabilities, as well as the observed prices, following Rust (1987). In particular, for a given set of parameters, we solve for the value functions, compute the implied conditional choice probabilities, calculate the likelihood and update the parameters until the latter is maximized (we calibrate the discount factor to $\beta = 0.995$).

Distribution of Freight Values and Freight Entry Costs For each observed price, we can back out the implied valuation v, using the price function, (3.5). In particular, note that we now have estimates for the freight matching probability, λ^f (it is given by m/f and f is estimated along with the matching function as described above) and the ship values, W_{ij} and J_i since ship costs are now known. Using data on the average value of trade in commodities and the average shipping price, we are able to estimate the bargaining power coefficient, γ, as well (the freight survival probability is calibrated to $\delta = 0.99$). We can now solve for v point-wise from the price equation and obtain the distributions $F_{ij}^v(.)$.

Finally, we recover the production and exporting costs, κ_{ij}, from the observed trade flows. Indeed, the probability that a potential exporter in i chooses to export in j is given by

$$G_{ij} = \frac{\exp\left(J_{ij}^f - \kappa_{ij}\right)}{1 + \sum_{l \neq i} \exp\left(J_{ij}^f - \kappa_{il}\right)} \tag{3.7}$$

China, Japan and Korea. We have also experimented with other groupings, such as (i) $c_{ij}^s = c^s$ for all i and j; (ii) coarser and finer groups; (iii) c_{ij}^s clustered by the distance between i and j, as well as clustered by the weather between i and j (both capture nonlinear effects of distance on the sailing cost).

where the values J_{ij}^f are now known from equations (3.3) and (3.4). Using our data to compute the number of loaded trips from i to j as well as external data on the total production of relevant commodities for each region i, we can estimate G_{ij} for all j, as well as G_{i0}. We recover κ_{ij} from (3.7) as in Berry (1994).

Results The main parameter estimates are given in Table 3.1. Briefly, we find that the fuel cost c_{ij}^s exhibits relatively low variation over space, consistent with the fact that fuel price dispersion is low. Port costs are large and vary substantially across regions. Freight valuations also differ significantly across regions: Brazil and then Australia are the highest value exporters. We also find that freight valuations correlate with the commodities produced in each region; for instance, countries that produce a lot of grain, which is the highest value commodity, tend to have higher valuations. Finally, the estimated matching function and freight distribution are reasonable: we find a high correlation between the retrieved number of the searching exporters and the realized exports (measured via Comtrade data). For further details on the parameter estimates, see BKP.

	Port Costs c_u	Sailing Costs c_s	Exporters Valuations μ_v	Preference Shock σ
North America West Coast	2.458	0.693	79.605	
	(0.07)	(0.002)	(2.038)	
North America East Coast	2.271	-	103.145	
	(0.021)	-	(2.229)	
Central America	1.846	0.693	73.161	
	(0.022)	(0.002)	(3.007)	
South America West Coast	1.996	0.693	59.063	
	(0.026)	(0.002)	(1.679)	
South America East Coast	2.563	-	125.877	
	(0.027)	-	(3.001)	
West Africa	1.421	0.64	57.838	
	(0.015)	(0.002)	(2.658)	
Mediterranean	1.637	0.568	59.87	
	(0.018)	(0.003)	(2.475)	
North Europe	1.399	0.568	49.199	
	(0.009)	(0.003)	(1.959)	
South Africa	2.478	0.64	99.074	
	(0.035)	(0.002)	(2.907)	
Middle East	1.273	0.568	44.203	
	(0.007)	(0.003)	(2.355)	
India	1.48	0.624	84.722	
	(0.014)	(0.003)	(4.2)	
South East Asia	1.67	0.56	72.282	
	(0.008)	(0.002)	(3.324)	
China	1.438	0.558	66.382	
	(0.01)	(0.002)	(3.61)	
Australia	2.635	0.56	70.507	
	(0.025)	(0.002)	(2.543)	
Japan-Korea	1.53	0.558	55.589	
	(0.022)	(0.002)	(2.514)	
				0.117
				(0.0008)

Note: all the estimates are in 100,000 USD. Standard errors computed from 500 bootstrap samples.

Table 3.1: Ship costs and exporter valuation estimates. The sailing cost for the East Coast of North and South America is set equal to 0.69 (the fuel cost).

3.4 Fuel Cost and Trade Elasticity

We now consider how a change in the cost of fuel affects maritime trade. We can think of this as the impact of an oil price shock on trade through transport costs. The overall effect is shown in Figure 3.1. The left panel plots the percent change in world exports, defined as the sum of realized matches between freights and ships over all ports, against the percent change in fuel costs captured by c_{ij}^s in our setup. The right panel plots the corresponding elasticity. The change in trade is substantial: the elasticity is estimated at 0.35 at the average fuel cost in our data and ranges from 0.1 to 1.2. This elasticity is lower than the results of the literature summarized in the Introduction.

A striking feature of the estimated trade elasticity, is its pronounced asymmetry in low and high levels of fuel costs. Indeed, the elasticity gets steeper as the fuel cost increases, while it plateaus at low fuel cost levels. For instance, consider a symmetric decline and increase of 30% from the weekly fuel expenditure in our sample of about 69,100$. When the fuel cost increases by 30%, total trade falls by about 15%; in contrast, when the fuel cost declines by 30%, total trade grows by about 8%, which is substantially less. As we argue below, this asymmetry is generated by the equilibrium of the transportation sector and in particular, the relative bargaining position of ships and exporters.

Since fuel costs are a core input in the transportation "production", higher fuel costs naturally lead to higher shipping prices, all else equal. Fuel costs, however, do not act just as inputs for ships; they also determine their relative bargaining position. When fuel costs increase, ships become "captive" to their current location: as it is costly to ballast elsewhere, ships are constrained in their ability to exploit opportunities in other regions. This limits the competition for ships between exporters across different regions and allows exporters to pay lower shipping prices than otherwise. In contrast, when fuel costs decrease, it is cheaper for ships to reallocate in space. Ships

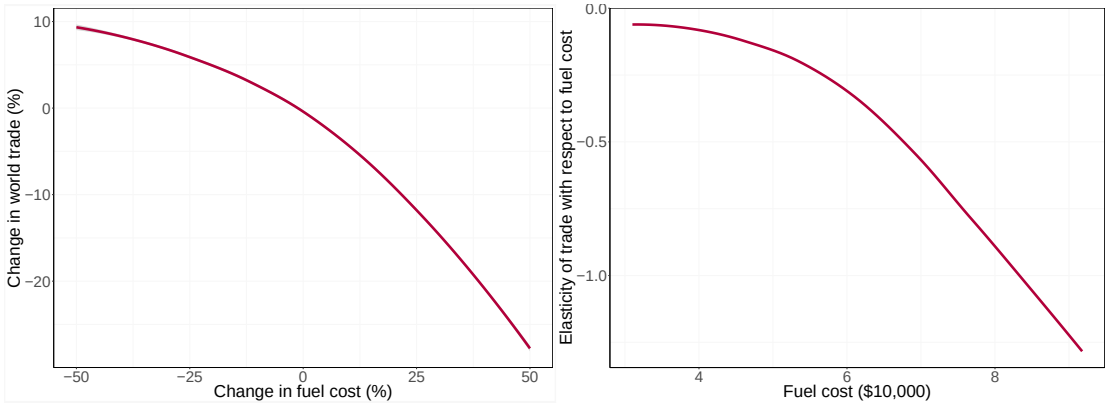

Figure 3.1: This figure uses the estimated model to compute changes in total trade, defined as the sum of realized matches of freights and ships over all ports, in reaction to changes in the fuel cost. The left panel shows the total change in world trade for different shocks to fuel costs, while the right panel plots the corresponding elasticity. As described in Section 3.3.3, in the baseline estimation we calibrate the c_{ij}^s to be equal to the average weekly fuel expenditure within our sample period, which is around \$69,100. To produce the figure above, we use the model to simulate the equilibrium level of exports associated with different counterfactual values of fuel expenditure c_{ij}^s. The range of counterfactual costs that we consider is from \$26,000 to \$80,000, which corresponds to plus and minus 50% changes to our baseline cost.

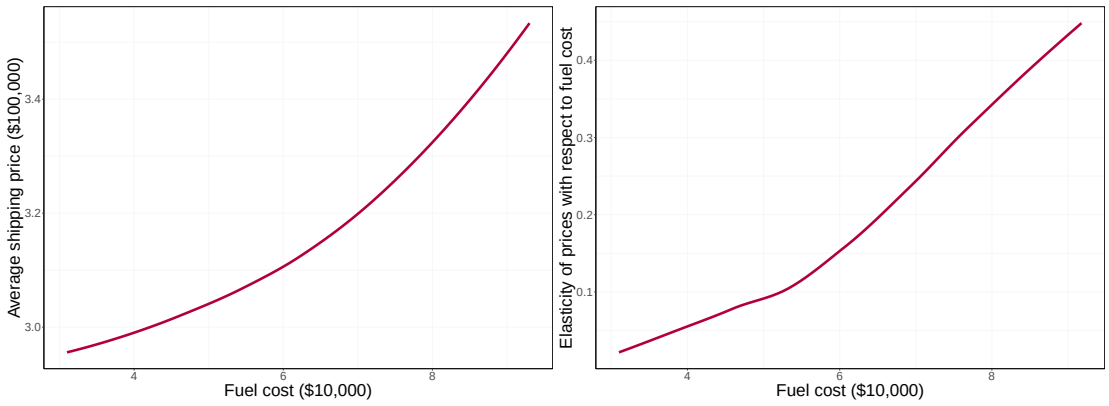

Figure 3.2: This figure uses the estimated model to compute shipping prices as a function of the fuel price. The left panel shows the average shipping prices for different fuel costs, while the right panel plots the corresponding elasticity. As described in Section 3.3.3, in our baseline estimation we calibrate c_{ij}^s to be equal to the average weekly fuel expenditure within our sample period, which is around \$69,100. To produce the figure above, we use our model to simulate the equilibrium value of shipping prices (τ_{ij}) associated with different counterfactual values of fuel expenditure c_{ij}^s. The range of counterfactual costs that we consider is from \$26,000 to \$80,000, which correspond to plus and minus 50% changes to our baseline cost.

are less "tied" to their current region and it is now the exporters that are in a weaker bargaining position and thus forced to agree to higher prices than otherwise. This indirect effect of fuel price shocks, coupled with the direct effect of fuel cost on the cost of transporting a cargo, determines the level and shape of the elasticity plotted in Figure 3.1.

Formally, from the price equation (3.5), a decline in the fuel cost c^s_{ij}, directly reduces the value of a traveling ship, W_{ij}, by reducing its cost (see equation (3.1)). At the same time however it affects ships' outside options, captured by the J_i term (see equation (3.3)). This effect tends to "dampen" the overall reduction in prices, and freight entry, when fuel costs fall.

In order to understand the asymmetry in Figure 3.1 note that the effect of fuel cost shocks on the ships' bargaining position is different for different levels of the fuel cost, c^s_{ij}. In a world with high fuel costs, ships' bargaining position reacts less to changes in the fuel cost. Indeed, at high fuel costs, the ship is likely to stay put rather than ballast, especially when in loading regions (net exporters). As a result, the price decline is not dampened as above and the increase in trade is sharper. Formally, when fuel costs are high, the outside option of ships J_i, is roughly equal to value of remaining unmatched in the current region, U_i, especially in exporting countries. A change in c^s_{ij} has little direct effect on U_i and therefore J_i, and the dampening impact of c^s_{ij} on prices is very small: most of the reduction in c^s_{ij} shows up directly in prices which leads to a correspondingly large increase in freight entry in net exporting countries.

Consider now a world with low fuel costs. In contrast to the situation above, it is now relatively cheap for ships to reallocate in space. In this "flat" world, a ship is likely to choose the ballasting rather than the waiting option, and as a result, declines in the fuel cost have a large impact on its outside option. This mutes substantially the increase in trade and the elasticity of trade shrinks dramatically. Formally, the

outside option of ships, J_i, is more likely to equal to one of the traveler values W_{ij} and a reduction in c^s_{ij}, can now have a sizable increase in , J_i , leading to a dampening in the price decline and a smaller increase in freight entry. It is worth noting that the same mechanisms are also present in an economy without search frictions, since ships' outside options are affected in a similar manner.

We also examine the pass-through of fuel cost shocks to shipping prices. Figure 3.2 displays shipping prices as a function of fuel costs, as well as the elasticity of shipping prices with respect to fuel costs. The elasticity is equal to 0.17 at the average fuel cost level and ranges from 0.03 to 0.43. Naturally, it features the same asymmetry. The level of the elasticity suggests that the pass-through of fuel costs to exporters is relatively low. Our estimates are similar to those found in Hummels (2007). The low pass-through, as well as the varying range of the elasticity suggests that transport costs are not well approximated by fuel costs and that modeling the equilibrium of the transportation sector is important in understanding the nature of trade costs.

We explore further our proposed mechanism by investigating separately the behavior of net exporters vs. net importers. Figures 3.3 and 3.4 plot exports, shipping prices and trade elasticities for different fuel costs levels for net exporting and net importing regions separately. The graphs are telling: at high levels of the fuel cost, for net exporters decreases in fuel cost are associated with substantially steeper increases in trade than for net importers. Indeed, ships in net exporting countries almost always prefer to stay put rather than ballast away when the fuel costs are high. Therefore, they benefit less from decreases in fuel prices – their bargaining position remains weak and any decline in fuel costs leads to large reductions in shipping prices and thus large increases in trade. Therefore, a decline in fuel cost when the fuel cost is high, disproportionally benefits net exporters, widening the gap between countries. As fuel costs further decline, however, the world becomes flat and importers and exporters become more similar to each other, as ships can costlessly reallocate.

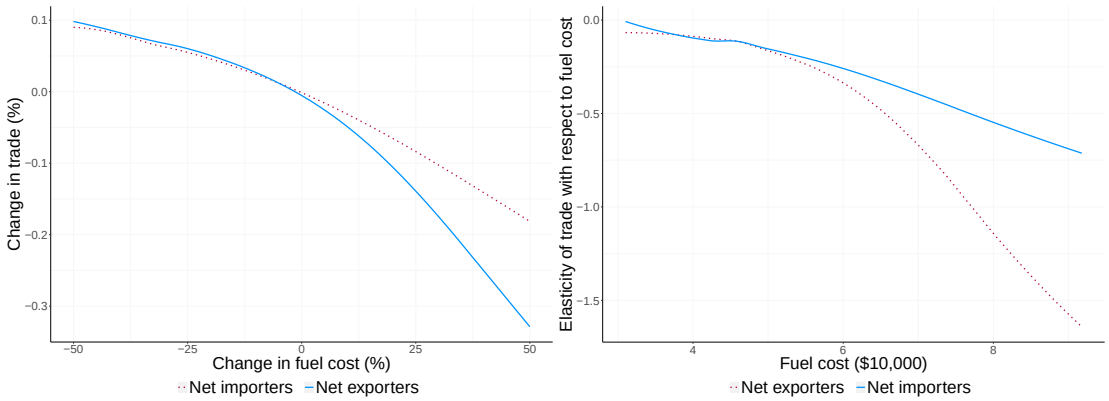

Figure 3.3: This figure uses the estimated model to compute changes in trade, defined as the sum of realized matches of freights and ships over all ports, in reaction to changes in fuel price for net importers and net exporters. The left panel shows the total change in trade for different shocks to fuel costs, while the right panel plots the corresponding elasticity. As described in Section 3.3.3, in our baseline estimation we calibrate c_{ij}^s to be equal to the average weekly fuel expenditure within our sample period, which is around \$69,100. To produce the figure above, we use our model to simulate the equilibrium values of trade associated with different counterfactual values of fuel expenditure c_{ij}^s. The range of counterfactual costs that we consider is from \$26,000 to \$80,000, which correspond to plus and minus 50% changes to our baseline cost. We classify a region as a net importer if the number of incoming cargoes is higher than the number of outgoing cargoes and vice versa.

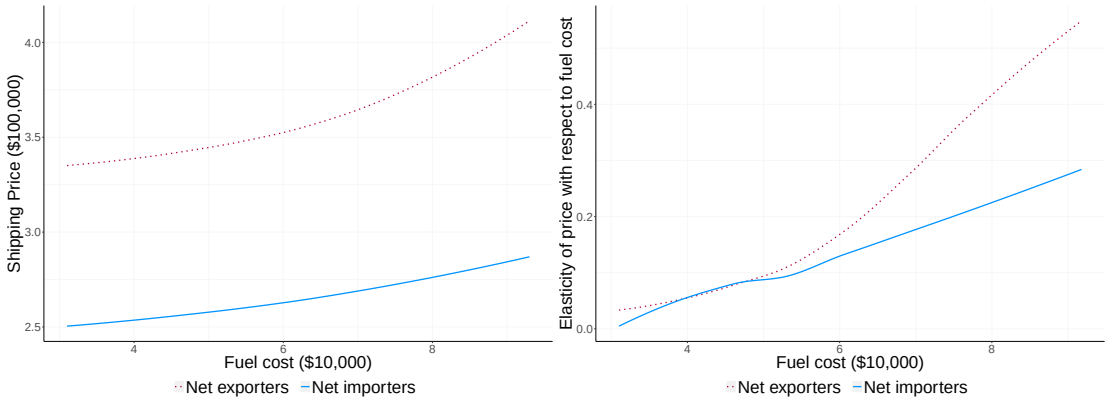

Figure 3.4: This figure uses the estimated model to compute shipping prices as a function of the fuel price for net importers and net exporters. The left panel shows the average shipping prices for different shocks to fuel costs, while the right panel plots the corresponding elasticity. As described in Section 3.3.3, in our baseline estimation we calibrate c_{ij}^s to be equal to the average weekly fuel expenditure within our sample period, which is around \$69,100. To produce the figure above, we use our model to simulate the equilibrium value of shipping prices (τ_{ij}) associated with different counterfactual values of fuel expenditure c_{ij}^s. The range of counterfactual costs that we consider is from \$26,000 to \$80,000, which correspond to plus and minus 50% changes to our baseline cost. We classify a region as a net importer if the number of incoming cargoes is higher than the number of outgoing cargoes and vice versa.

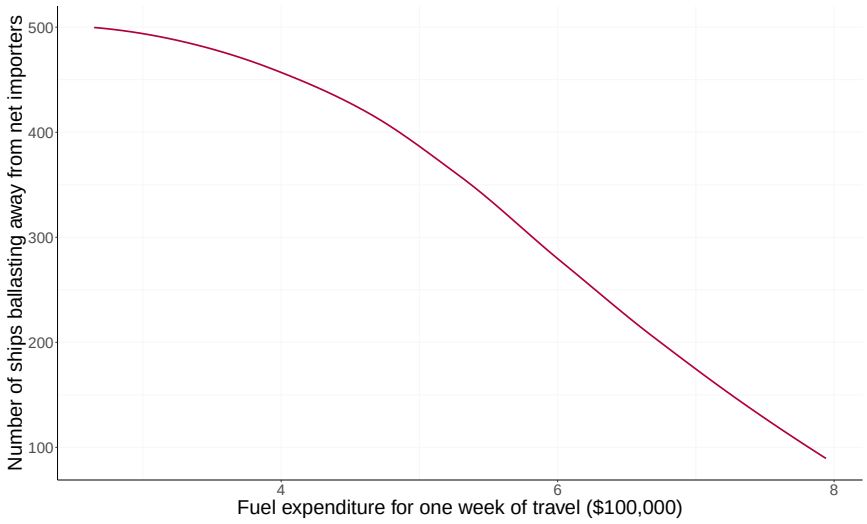

Figure 3.5: This figure uses the estimated model to compute the number of ships ballasting away from net importers as a function of the fuel price. As described in Section 3.3.3, in our baseline estimation we calibrate c_{ij}^s to be equal to the average weekly fuel expenditure within our sample period, which is around \$69,100. To produce the figure above, we use our model to simulate the equilibrium number of ballasting ships associated with different counterfactual values of fuel expenditure c_{ij}^s. The range of counterfactual costs that we consider is from \$26,000 to \$80,000, which correspond to plus and minus 50% changes to our baseline cost We classify a region as a net importer if the number of incoming cargoes is higher than the number of outgoing cargoes.

Finally, it is worth noting that, as fuel costs decline ships do ballast more and there is a reallocation across space from net importers to net exporters. However at low levels of fuel costs, further declines lead to small changes in reallocation: at these levels most opportunities are taken up and there are considerably fewer ships waiting at net importers that are able to take advantage of the reduction in fuel cost to ballast to net exporters. This is shown in Figure 3.5 where the number of ships ballasting away from net importers flattens out at sufficiently low levels of the fuel cost. This also implies a much lower trade elasticity to oil shocks, as shown in Figure 3.1. The world fleet utilization is close to full capacity, given the geographical constraints.

3.5 Oil Shocks and Price Dispersion

In this section we search for descriptive evidence that the equilibrium in the transportation sector and the relative bargaining position of ships and exporters are important determinants of world trade elasticities with respect to oil shocks. To do so, we consider a prediction of our setup regarding the correlation between fuel oil prices and the cross-sectional dispersion of shipping prices.

Indeed, consistent with the narrative of the previous section, as fuel costs decline and the world becomes flat, ship values tend to equalize over space. For instance, a ship in China is now not much worse off than a ship in Australia, as ballasting from China to Australia is cheaper. As the dispersion in ship value functions declines, so does the dispersion in shipping prices

To make this argument more precise, consider again the price equation, (3.5), for a given origin i: differences in prices across different destinations j are driven by differences in the value function of a traveling ship, W_{ij} (holding constant the distribution of freight valuations). As c^s_{ij} falls, W_{ij} tend to equalize across *destinations*, j, both because differences in distance are less important, but also because differences in the

ship's value, U_j, across different destinations matter less. As discussed above, ending a trip in a destination that is attractive to ships is less important than it used to be due to the cheaper ballasting.

In addition, again from the price equation, (3.5), holding constant differences in freight valuations and matching probabilities, any difference in the shipping price across *origins*, i, is driven by differences in the value of an unmatched ship, J_i.[8] From the discussion above, as fuel costs fall, so do the differences in J_i across i. Nash bargaining implies that all else equal, differences in prices should also shrink.

It is worth emphasizing that the decline in the dispersion of prices *per-day* both across destinations, but also across origins, is driven by the convergence in the ship valuations across space, resulting from a lower cost of ballasting. In a model where prices depended exclusively on distances and ships' outside options did not affect them, the dispersion of prices *per-day* should not depend on the level of fuel costs.[9]

Therefore, a testable implication of our framework is that shipping price dispersion should be increasing the fuel cost. Tables 3.2 and 3.3 report the results from a regression of the dispersion of shipping prices on fuel cost. Table 3.2 shows that as fuel prices increase, the dispersion of per-day shipping prices for trips with different origins increases. Similarly, Table 3.3 shows that as fuel prices increase, the dispersion of shipping prices per-day for trips with different destinations increases. In both cases, the coefficient of the fuel cost is positive and significant, and robust to the inclusion of different time fixed effects.

[8]Note that, as shown in equation (3.1), conditions at the origin do not affect the value of W_{ij}.

[9]It should be noted that search frictions are not necessary for this result. In a world without search frictions, but where prices are still endogenous, ships' outside options continue to matter and are affected in a similar manner.

	I	II	III
	Dispersion of per-day-shipping price across different origins, log		
Fuel Price (log)	0.288^{***}	0.741^{***}	0.597^{**}
	(0.090)	(0.194)	(0.236)
FE	None	Year	Quarter × year
Observations	175	175	175
R^2	0.056	0.735	0.917

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

Table 3.2: This table reports the estimates from a regression of the dispersion of shipping prices across trip origins on fuel costs. For all the regressions, we compute the average shipping price per-day that different exporting countries face within a month. The dependent variable in all regressions is the standard deviation of this average shipping price across origins (in logs). The main independent variable is the monthly fuel price (in logs). In column I we report the raw correlation, while in columns II and III we add year and quarter-year fixed effects respectively.

3.6 Fleet Fuel Efficiency and World Trade

In this section, we compute how much recent trends in the fleet's fuel efficiency have affected trade flows. Figure 3.6, which is taken from Faber and Hoen (2015), plots an index of energy efficiency over the last 50 years, for different types of ships.[10] Energy efficiency for bulk carriers improved dramatically (by about 25%) in the 1980s and

[10]In order to assess design fuel efficiency, Faber and Hoen (2015) use data from the IHS Maritime World Register of Ships and from Clarksons World Fleet Register to compute each ship's Estimated Index Value (EIV) and compare it with the EIVs of ships that entered the fleet between 1999 and 2008 (reference line). A value above the reference line means that a ship emits more CO2 per ton-mile under standard conditions, and it is therefore less efficient than the average comparable ship between 1999 and 2008. The EIV computation takes into account the ship's capacity (deadweight tonnage), main engine power, auxiliary power and design speed. According to the authors, the EIV is a simplified version of the EEDI shown in Figure 3.6. The interested reader should refer to Faber and Hoen (2015) for more details.

after a short stable period began deteriorating. This observed reversion in energy efficiency is striking and not noted in other transport sectors, such as aircraft or trucks.

	I	II	III
	Dispersion of per-day-shipping price across different destinations, log		
Fuel Price (log)	0.228^{**}	0.812^{***}	0.766^{***}
	(0.091)	(0.200)	(0.229)
FE	None	Year	Quarter × year
Observations	175	175	175
R^2	0.035	0.718	0.923

$^{**}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

Table 3.3: This table reports the estimates from a regression of the dispersion of shipping prices across trip destinations on fuel costs. For all the regressions, we compute the average shipping price per-day that different exporting countries face within a month. The dependent variable in all regressions is the standard deviation of this average shipping price across destinations (in logs). The main independent variable is the monthly fuel price (in logs). In column I we report the raw correlation, while in columns II and III we add year and quarter-year fixed effects respectively.

Although it may seem unintuitive at first, the reversion makes sense if one considers the determinants of design fuel efficiency. Indeed, ship design is determined endogenously by the "long-run" market equilibrium in the shipping and shipbuilding markets; as such it is affected by market conditions (predominantly world trade), fuel costs, as well as environmental policies. For instance, the sharp improvements in efficiency in the 1980s follow the oil crisis of the 1970s. The lag in the shipyards' reaction is also consistent with the time required to produce novel ship designs (Faber and Hoen (2015)). The current deterioration in ship efficiency may originally have been due to the massive increase in trade starting in the 1990s-2000s: at the

time, shipyards faced severe capacity constraints and opted for simple and quick to build designs (Kalouptsidi (2014), Kalouptsidi (2017), Faber and Hoen (2015)), while high freight rates made high fuel costs less painful for shipowners. Nowadays, the persistently low fuel costs reduce the shipowners' willingness-to-pay for high energy efficiency.

Overall, the combination of (i) long-term trends in the shipbuilding industry; (ii) technological improvements; and (iii) environmental policies lead to time-varying ship designs in terms of fuel efficiency. Here, we use our estimates for the trade elasticity to calculate (i) the gains in trade because of the improvement in design fuel efficiency since the early 1980s; and (ii) the reduction in trade brought about by the more recent deterioration in fuel efficiency. Our estimated elasticities suggest that the 25% improvement in fuel efficiency since the 1980s have led to a decline in shipping prices by 5.5% and a corresponding increase in trade by 11.8%. On the other hand, since the 1990s efficiency has deteriorated by 13%, resulting in a 5.6% reduction in world trade, and a 2.5% increase in shipping prices.

In summary, consistent with the trade elasticity estimated above, we find that ship design is an important determinant of trade. As global institutions are currently considering the phasing in of substantial environmental policies (e.g. the International Maritime Organization (IMO) is currently phasing in limits on sulphur in ship fuel) it is our hope that our methodology and estimates can be of use in the cost-benefit analysis of determining the optimal levels of environmental standards.

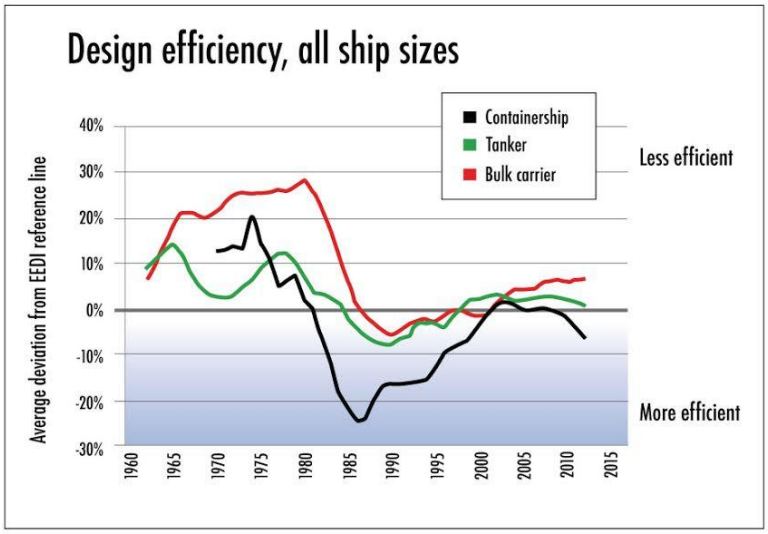

Figure 3.6: This figure displays the fleet's Energy Efficiency Design Index (EEDI) compared to the EEDIs of ships that entered the fleet between 1999 and 2008 (reference line). The graph is taken from Faber and Hoen (2015). See also footnote 10.

3.7 Conclusions

In this paper, we measure the world trade elasticity with respect to oil prices. We show that the elasticity is substantially asymmetric with respect to high and low fuel costs. In particular, as fuel costs decline, the elasticity becomes flatter. This "flattening-out" of the elasticity is attributed to the stronger bargaining position of ships vs. exporters as distances becomes less important. Indeed, as fuel costs decline and the world becomes flat, ships can leverage their ability to reallocate cheaply and command higher prices. Their stronger bargaining position puts a break on trade growth. The trade elasticity with respect to fuel costs is of great relevance in a number of policy debates regarding environmental, trade and other regulations.

Our approach is not free of caveats. We do not model all the margins along which ships can react to fuel costs. For instance, ships can adjust their speed in response to oil shocks and in particular they can "slow-steam" during oil price spikes. Moreover, we only consider our model in steady state; therefore, firms do not form expectations about oil price fluctuations, nor do they act dynamically with respect to such beliefs. Our elasticity should be considered a short or medium-term one, as we do not model the investment decisions of shipowners in purchasing vessels, nor the reaction of the shipbuilding sector that may change the design of ships. All these issues form interesting avenues for future research. Finally, we consider the impact of oil shocks solely through transportation costs. This is the stated goal of the paper; yet, oil is an key input in production of several industries and understanding its overall impact in the economy is an important question of interest.

References

Adland, R., Fonnes, G., Jia, H., Lampe, O. D. and Strandenes, S. P.: 2017, The impact of regional environmental regulations on empirical vessel speeds,

Transportation Research Part D **53**, 37–49.

Adland, R. and Jia, H.: 2016a, Dynamic speed choice in bulk shipping, *Maritime Economics & Logistics* pp. 1–14.

Adland, R. and Jia, H.: 2016b, Vessel speed analytics using satellite-based ship position data, *Proceedings of the 2016 IEEE IEEM* pp. 1299–1303.

Anderson, J. E. and Van Wincoop, E.: 2003, Gravity with gravitas: A solution to the border puzzle, *American Economic Review* **93**(1), 170–192.

Asturias, J.: 2016, Endogenous transportation costs, *mimeo, School of Foreign Service in Qatar, Georgetown University* .

Backus, D. K. and Crucini, M. J.: 2000, Oil prices and the terms of trade, *Journal of International Economics* **50**(1), 185–213.

Berry, S.: 1994, Estimating discrete-choice models of product differentiation, *RAND Journal of Economics* **25**(2), 242–262.

Brancaccio, G., Kalouptsidi, M. and Papageorgiou, T.: 2017, Geography, search frictions and endogenous trade costs, *NBER working paper* **23581**.

Bridgman, B.: 2008, Energy prices and the expansion of world trade, *Review of Economic Dynamics* **11**(4), 904–916.

Buchholz, N.: 2016, Spatial equilibrium, search frictions and efficient regulation in the taxi industry, *mimeo, Princeton University* .

Chen, S.-S. and Hsu, K.-W.: 2012, Reverse globalization: Does high oil price volatility discourage international trade, *Energy Economics* **34**(5), 1634–1643.

Ericson, R. and Pakes, A.: 1995, Markov-perfect industry dynamics: A framework for empirical work, *The Review of Economic Studies* **62**(1), 53–82.

Estevadeordal, A., Frantz, B. and Taylor, A. M.: 2003, The rise and fall of world trade, 1870-1939, *Quarterly Journal of Economics* **118**(2), 359–407.

Faber, J. and Hoen, M.: 2015, Historical trends in ship design efficiency, *Technical report*, CE Delft and Transport & Environment, Delft, Netherlands.

Frechette, G. R., Lizzeri, A. and Salz, T.: 2016, Frictions in a competitive, regulated market evidence from taxis, *mimeo, New York University* .

Hopenhayn, H. A.: 1992, Entry, exit, and firm dynamics in long run equilibrium, *Econometrica* **60**(5), 1127–1150.

Hummels, D.: 2007, Transportation costs and international trade in the second era of globalization, *The Journal of Economic Perspectives* **21**(3), 131–154.

Kalouptsidi, M.: 2014, Time to build and fluctuations in bulk shipping, *American Economic Review* **104**(2), 564–608.

Kalouptsidi, M.: 2017, Detection and impact of industrial subsidies, the case of chinese shipbuilding, *forthcoming, Review of Economic Studies* .

Kilian, L., Rebucci, A. and Spatafora, N.: 2009, Oil shocks and external balances, *Journal of International Economics* **77**(2), 181–194.

Krugman, P.: 1991, Increasing returns and economic geography, *The Journal of Political Economy* **99**(3), 483–499.

Lagos, R.: 2000, An alternative approach to search frictions, *The Journal of Political Economy* **108**(5), 851–873.

Limao, N. and Venables, A. J.: 2001, Infrastructure, geographical disadvantage, transport costs, and trade, *The World Bank Economic Review* **15**(1), 451–479.

Matzkin, R. L.: 2003, Nonparametric estimation of nonadditive random functions, *Econometrica* **71**(5), 1339–1375.

Mortensen, D. T. and Pissarides, C. A.: 1994, Job creation and job destruction in the theory of unemployment, *The Review of Economic Studies* **61**(3), 397–415.

North, D.: 1958, Ocean freight rates and economic development 1750-1913, *The Journal of Economic History* **18**(4), 537–555.

Petrongolo, B. and Pissarides, C. A.: 2001, Looking into the black box: A survey of the matching function, *Journal of Economic Literature* **39**(2), 390–431.

Rubin, J.: 2009, *Why Your World Is About to Get a Whole Lot Smaller: Oil and the End of Globalization*, Random House, New York.

Rust, J.: 1987, Optimal replacement of gmc bus engines: An empirical model of harold zurcher, *Econometrica* **55**(5), 999–1033.

Stopford, M.: 2009, *Maritime Economics*, 3rd edn, Routledge, London and New York.

UNCTAD: 2015, Review of maritime transport, *Technical report*, United Nations Conference on Trade and Development, Geneva, Switzerland.

von Below, D. and Vezina, P.-L.: 2016, The trade consequences of pricey oil, *IMF Economic Review* **64**(2), 303–318.

Wong, W. F.: 2017, The round trip effect: Endogenous transport costs and international trade, *mimeo,University of Wisconsin-Madison* .

Wood, J.: 2008, Fuelling global trade: How gdp growth and oil prices affect international trade flows, *Technical report*, The Economist Intelligence Unit, London, New York, Hong Kong.