



City Research Online

City, University of London Institutional Repository

Citation: Medina, Anthony (2021). A Common Value OTC Network Model. (Unpublished Doctoral thesis, City, University of London)

This is the accepted version of the paper.

This version of the publication may differ from the final published version.

Permanent repository link: <https://openaccess.city.ac.uk/id/eprint/27310/>

Link to published version:

Copyright: City Research Online aims to make research outputs of City, University of London available to a wider audience. Copyright and Moral Rights remain with the author(s) and/or copyright holders. URLs from City Research Online may be freely distributed and linked to.

Reuse: Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

CITY, UNIVERSITY OF LONDON

DOCTORAL THESIS

A Common Value OTC Network Model

Author:

Anthony MEDINA

Supervisor:

Prof. Giulia IORI

*A thesis submitted in fulfilment of the requirements
for the degree of Doctor of Philosophy
in the*

Department of Economics

August 31, 2021

Declaration of Authorship

I, Anthony MEDINA, declare that this thesis titled, A Common Value OTC Network Model and the work presented in it are my own. I confirm that:

- This work was done wholly or mainly while in candidature for a research degree at this University.
- Where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated.
- Where I have consulted the published work of others, this is always clearly attributed.
- Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work.
- I have acknowledged all main sources of help.
- Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself.

Signed: 

Date: 31 August 2021

There are known knowns. These are things we know that we know. There are known unknowns. That is to say, there are things that we know we don't know. But there are also unknown unknowns. There are things we don't know we don't know. ... it is the latter category that tend to be the difficult ones.

– Donald Henry Rumsfeld, US Secretary of Defence, (2002)

Abstract

Anthony MEDINA

A Common Value OTC Network Model

This thesis looks at decentralized client-dealer markets by representing them as a Bertrand competition over a network of connections and examines the Nash equilibrium bid-ask spread properties. We specifically look at OTC financial markets with a common but unknown true value and show how equilibrium bid and ask prices can be viewed as the optimal amounts of signal reduction in a common-value first-price auction. We algebraically examine the equilibrium in the duopoly case and then analyze the equilibriums numerically in the n-dealer case. Our model suggests that changing transparency can have non-intuitive effects that depend on both the network structure and the information asymmetry. Similarly, the model can explain the empirically observed centrality premiums and discounts found in some OTC markets and other noted pricing anomalies. We apply our network model to the UK betting market using data from the UK General Election 2019. In betting markets, win probabilities are the asset, and dealers are the bookmakers, and we find that bookmakers' odds closely follow the model's predicted levels and exhibit a predicted centrality discount.

Acknowledgements

Many thanks to my Supervisor, Giulia Iori, for her endless patience and guidance. Thanks also to my family, who have been a constant source of both motivation and pleasant distraction.

Contents

Declaration of Authorship	iii
Abstract	vii
Acknowledgements	ix
1 Introduction	1
1.1 General	1
1.2 Related Literature	5
1.2.1 OTC Literature	5
1.2.2 Microstructure Literature	7
Inventory Models	8
Asymmetric Information Models	9
2 A Model of Trading in OTC Markets	13
2.1 Agent Types	13
2.1.1 Clients	13
2.1.2 Dealers	14
2.2 Common Asset True Value and Signals	14
2.3 Client and Dealer Strategies	17
2.4 Connectivity in a OTC Financial Market	18
2.5 OTC Trading Architecture	19
2.6 Model Protocols	20
2.7 Payoffs	21
2.7.1 N dealers, 1 Client payoffs	22
2.8 Bid-Ask Spreads, Bertrand Competition and Auctions	26
2.8.1 Bid-Ask Spreads and Bertrand Competition	26

2.8.2	Bertrand Competition and Auctions	27
2.9	Equilibrium Bidding Strategies in First Price Common Value Auctions	28
2.9.1	Introduction	28
2.9.2	Setup	28
2.9.3	Solution to the Standard FPSB Auction	29
2.9.4	Equivalence of Reservation Prices and Extra Bidder	30
2.9.5	Signal and Bid Distributions	31
2.9.6	Equilibrium Bidding Strategy	32
2.9.7	Summary	34
2.9.8	Numerical Checks	35
3	Duopoly OTC Markets	37
3.1	Introduction	37
3.2	The Model	40
3.2.1	Information Sets	41
3.2.2	Normal Distribution of the Estimation (Signal) Errors	41
3.3	Optimal Bidding Strategies	42
3.3.1	The Standard Auction Model	42
	The Wilson Model	42
3.3.2	Linear Bidding Functions	44
3.4	Standard Duopoly Market	45
3.4.1	Assumptions	45
3.4.2	Game setup	46
3.4.3	Expected Payoffs	46
3.4.4	Equilibrium	48
	Symmetric Information Case	48
	Asymmetric Information	49
3.4.5	Customer Payoffs	53
3.4.6	Non Unit Demand Functions	54
3.5	OTC Market Bid-Ask Spreads	56
3.6	Client Reservation Prices	57
3.7	Other Informational Asymmetries	59

3.7.1	Client-Dealer Asymmetry in Monopoly	59
3.7.2	Client-Dealer Asymmetry in Duopoly	60
3.7.3	Price Dispersion	61
3.8	Conclusion	63
4	OTC Network Model	65
4.1	Introduction	65
4.2	Basic Network and Graph Concepts	66
4.2.1	Basic Definitions	66
4.2.2	Types of Networks	67
4.2.3	Centrality Measures	68
4.2.4	Random Networks	70
4.2.5	Core-Periphery Networks	71
4.2.6	Degree Distributions	71
4.3	Empirically Observed Market Networks	73
4.4	Preferential Attachment Network Models	74
4.4.1	Bipartite Model with Preferential Selection	74
4.4.2	Higher Degree Client Node Networks	76
4.5	N Dealer, M Client Network Game	78
4.6	Network Payoffs	80
4.7	Weighted Biadjacency Matrix Representation	81
	Monopoly case	82
	Duopoly Case	83
	General case	83
4.7.1	Dealer Payoff in Weighted Network Representation	84
	Example in a 3-dealer network	85
4.8	Simplified Preferential Attachment Network	86
4.8.1	1 and n Degree Client Case	86
4.9	The OTC Network as a Static Bayesian Game	89
4.10	Equilibrium Pricing Solutions	92
4.10.1	Existence	92
4.10.2	Case 1 - the regular network	93

4.10.3	Case 2 - the irregular network	94
4.10.4	Numerical Solution of equilibrium	96
4.10.5	Algorithm Convergence to NE	98
5	Equilibrium Pricing in the Network	101
5.1	Introduction	101
5.2	Bid-Ask Spreads and Uncertainty	102
5.3	Bid Ask Spreads and Transparency	104
5.3.1	Preferential Selection Network Payoffs	105
5.3.2	The Regular Network	107
5.3.3	Bipartite Preferential Attachment Network	109
5.4	Centrality Premiums in the Network Model	112
5.4.1	Stylised Model of OTC Network	113
	Case 1 - Dealers Better Informed than Clients	114
	Case 2 - Dealers Equally (or Worse) Informed than Clients	116
5.4.2	1 and 2 Degree Client Case	118
5.5	Empirical Anomalies in OTC Markets	121
5.5.1	Low-Risk and High-Risk Bonds Bid-Ask Spreads	121
5.5.2	Liquid and Illiquid Bonds Bid-Ask Spreads	123
5.6	Conclusion and Discussion	123
6	Gambling Markets as Common Value OTC Markets	125
6.1	Introduction	125
6.2	Literature Review	128
6.3	Background	129
6.3.1	UK General Election Basics	129
6.3.2	Vote to Seat Prediction Problem	130
6.4	Probabilities and Bid-Ask Spreads from Betting Data	131
6.4.1	Dataset	131
6.4.2	Implied Probability Bid-Ask Calculation	132
6.4.3	Summary Win Probabilities	134
6.4.4	Bilateral and Exchange Pricing	135
6.4.5	Adjusting Implied Probabilities from Bias	136

6.5	Implied Beliefs and Opinion Polls	138
6.5.1	Implied Belief Calculation	138
6.5.2	Opinion Polling and Expected Beliefs	141
6.6	Model	143
6.6.1	Network Structure of the Betting Market	143
6.6.2	Multi-Period Unknown Common Value FPSBA Model	146
6.6.3	Converting Vote Share to Seat Majority Beliefs	150
6.7	Observed and Predicted Spreads	152
6.8	Model Evaluation	154
6.9	Centrality Premium Effect	156
6.10	Conclusion	159
7	Conclusion and Discussion	161
A	Appendix To Chapter 3	175
A.0.1	Derivation of Monopoly Dealer Payoff Function	175
B	Appendix To Chapter 4	177
B.0.1	Formal Specification of Static Bayesian Games	177
B.0.2	Convergence with an Agent Based Model (ABM)	178

List of Figures

2.1	Illustration of the normal approximation to lognormal when means are greater than standard deviations	16
2.2	A graphical bipartite network representation of an OTC market	19
2.3	The trading protocols for a buying client connected to a subset of dealers	21
2.4	Comparison of numerical simulation to the derived payoff functions . .	25
2.5	Sample of random realizations in each auction	26
2.6	Best response functions of the 2 strategic bidders	36
2.7	Optimal bid reductions for the n strategic bidders with 1 naive bidder wrt relative signal variance	36
3.1	Theoretical and Numerical Simulation of Signal Reduction Best Response Curves - Firm1(blue) $N(0,1)$, Firm2(red) $N(0,1)$ signal error.	49
3.2	Signal Reduction Best Response Curves - Firm 1 has $N(0,1)$ and Firm 2 has $N(0,2)$ signal error	50
3.3	Numerically Derived Signal Reduction Best Response Curves - Firm 1 has $N(0,1)$ and Firm 2 has $N(0,4)$ signal error with a (-)ve sloping linear demand curve	55
3.4	Payoff surfaces for the 2 firms with the best responses shown in red and blue	59
3.5	Monopoly dealer equilibrium spread levels varying with client relative signal variance	60
3.6	Dealer equilibrium spread level with client signal variance	61
4.1	Degree distribution of Poisson and Power Law Networks	72
4.2	Clients C_1, \dots, C_5 and associated neighbourhood sets and valuations . .	79

4.3	Example of a 5-dealer market with clients C1,C2, C3 and associated neighbourhood sets	81
4.4	Strategic equivalence of monopoly market, M clients	82
4.5	Strategic equivalence of duopoly market, M total clients	83
4.6	Relationship between α and weighted degree centrality with 50% low and high degree clients	89
4.7	Relationship between α and relative degree centrality $DC(D_n)/DC(D_p)$ of central to peripheral dealers with various proportions of low degree clients	89
4.8	Dealer payoff function wrt spread choice, keeping other dealer spreads fixed	96
4.9	Strategic equivalence of networks	99
4.10	Network with client grouping, $G(D_i, D_j)$, represents game G played with dealers D_i and D_j	99
4.11	Algorithm convergence to Nash equilibrium in the 2 networks, N(0,1) signal errors	100
5.1	Dealer spreads in US Treasuries	104
5.2	Dealer spreads to degree and asymmetric information	108
5.3	Effects of increasing information transparency with client degree when clients have high signal error	109
5.4	Dealer spreads to degree and asymmetric information	110
5.5	Equilibrium Spreads with Asymmetric Information	110
5.6	Centrality premium sensitivity to various network configurations . . .	114
5.7	Centrality premium sensitivity to increasing asymmetric information .	115
5.8	centrality discount when clients equally informed	116
5.9	Envelope of centrality premia with asymmetric information	117
5.10	Raw centrality of central dealer relative to changes in network parameters α, β	120
5.11	Relationship between α and relative degree centrality $DC(D_n)/DC(D_p)$ of core to peripheral dealers with various proportions of low degree clients	120

5.12 At the equilibrium, when clients are equally informed there is a centrality discount and when the clients are worse informed there is a centrality premium	121
6.1 Vote Share Thresholds to Win Election (electoralcalculas.com)	131
6.2 Sample daily odds prices from main bookmakers	132
6.3 Summary of Implied Bid Offer Probabilities	134
6.4 Average Implied Win Probabilities from Odds Prices	135
6.5 Observe Exchange, Bookmaker and Best of Bookmaker Spreads	136
6.6 Sample interval level bets from 888Sport bookmaker	138
6.7 Histograms of Implied Observed Belief Distributions	140
6.8 Implied Vote Shares	141
6.9 Implied SD of Votes Shares	141
6.10 Opinion Poll Forecasts	142
6.11 Market Forecasts	142
6.12 Average Number of Online Accounts, taken from the Gambling Commission	144
6.13 Illustration of Winners Curse due to Overrounds	149
6.14 Standard Deviation of Conservative Party Most Seats Win Belief Probabilities	151
6.15 Mean Observed Spreads and Predicted Spread Levels Evaluated from Belief Data	152
6.16 Observed and Belief Smoothed Predicted Spread Levels	153
6.17 Observed vs Predicted Linear Regression	154
6.18 Residuals of observed vs predicted values regression model	154
6.19 F-test results on the linear regression coefficients	155
6.20 Partitioning on Bookmakers by Revenue	157
6.21 Spreads of Largest and Smallest Bookmakers	158
B.1 convergence of dealers' equilibrium spreads in ABM	179

List of Tables

6.1	Win and Lose Probabilities derived for Odds Prices	133
6.2	Summary of Conservative Party Event Odds and Implied True Win Probabilities using N and P method	139
6.3	Average Test Statistics for Normality of Data	140
6.4	Opinion Poll Summary - Election Day -1	143

Chapter 1

Introduction

example.bib

1.1 General

Many standard financial market instruments, such as common stocks or listed options, trade on organized exchanges such as the NYSE or the CBOE. However, a significant number of financial assets, such as most loans, government and corporate bonds, derivatives, and currencies, have no central exchange, and traders need to locate a counterparty for each trade. The reason is that, unlike equities, many securities are not as homogenous, with non-standard expiries and payoffs. In the EU, for example, there are 6800 listed shares, but Xtrakter's CUPID database contains information on over 150,000 debt securities in issue, with only 3000 that trade more than once per day. This decentralization has led to professional firms becoming intermediaries or market-makers, and these markets are known as over-the-counter or OTC markets. In these markets, transactions are bilateral, prices are dispersed, trading relationships are persistent, and typically, a few large dealers intermediate a large share of the trading volume, [Babus and Kondor \(2018\)](#). These markets are also characterized by dealer sophistication due to the illiquid and sometimes complex nature of the instruments traded. This thesis develops a network model of OTC trade between clients and dealers that reflects these features.

Our model incorporates two critical drivers of price formation in OTC client-dealer markets. Firstly, the extent of the participants' information sets over the true asset value, a common theme in standard microstructure models, which we model with a noisy signal or estimate of the true value. Secondly, the topology of the network of

connections between clients and dealers, as seen in the large body of work that views markets as a network. We explore a network model that combines these two drivers, and we examine the Bayesian Nash Equilibrium dealer pricing strategies.

We investigate how increasing transparency using our network model, where transparency is modeled as a refining of the clients' information sets regarding the asset's true value. We find that the clients' degree distribution is critical in the direction of the effect on spreads. Clients with low dealer connectivity (a degree of 1 or 2) benefit from increased transparency through lower bid-ask spreads. In contrast, clients with a higher degree (3 or more links) suffer increased spreads in equilibrium. This increase is due to the increasing dominance of the winners' curse effect and the initial relative informedness of clients and dealers. This result might offer some explanation as to the mixed results when regulators have mandated increased transparency.

An effect in OTC markets, known as the centrality premium effect, where more centrally located dealers make consistently wider or tighter spreads than less centrally located dealers, has been documented by various authors. Traditional market models cannot totally explain the range of these effects. [Hollifield et al. \(2012\)](#) find a negative relationship between bid-ask spreads and dealer centrality in the US securitization market. In contrast, [Di Maggio et al. \(2017\)](#) and [Li and Schürhoff \(2019\)](#) find a positive relationship between bid-ask spreads and dealer centrality in the US corporate bond and US municipal bond market. Our network model predicts these premiums and discounts. Similarly, our model correctly predicts a centrality discount for the largest bookmakers in our empirical study of the UK gambling market.

We also apply the model to several pricing anomalies identified by [Biais et al. \(2005\)](#), such as why safe municipal bonds have bigger spreads than riskier bonds or equities. Our model explains these effects by the different link distributions of the clients and asymmetric information. The effect of illiquid bonds having tighter spreads than very liquid bonds in the US corporate bond market was observed by [Goldstein et al. \(2007\)](#). Our model shows how these two markets' topological and information structures are consistent with this effect.

In our model, the market consists of clients, a set of dealers, and a homogenous security that has a common but unknown value to all agents. It is a natural assumption to assume a common value element in security valuation, [Biais et al. \(2000\)](#).

Clients requests trading quotes from their connected dealers, as determined by their links in a network. Each client has an exogenous reason for trading but is not a noise trader in the traditional microstructure sense. Instead, they are boundedly rational and attempt to measure the true value before trading. The clients then trade at the best price observed in their local network, only if the price is better than their estimate. This classification of client behavior is consistent with the empirical work of [Gode and Sunder \(1993\)](#), where the machine traders that turned out to be most similar to human traders were the zero-intelligence ones with budget constraints - traders who were not allowed to trade at a perceived loss. Similar to [de Kamps et al. \(2012\)](#), we examine heterogeneous boundedly rational clients who base their trading decisions on their noisy valuation of the asset. Dealers are risk-neutral and strategic (they form expectations over other beliefs and actions in the network) and, similar to the clients, estimate the asset's true value. They use these to make a bid and ask price to maximize their payoffs in competition with the other connected dealers.

This thesis follows many authors in representing decentralized trade using a network, and we represent a stylized OTC market as a bipartite network, where node types are partitioned into clients and dealers. This partition follows naturally from price makers and price takers' market functions, leading to a Bertrand price competition among connected dealers, with the best-priced dealer winning the trade.

In contrast to the majority of papers that examine the OTC inter-dealer market, this thesis examines the client-dealer OTC network in a similar way to authors such as [Hendershott et al. \(2020\)](#), and [de Kamps et al. \(2012\)](#). Our interest is how clients connect to the dealer set and use the empirical results of Hendershott from the corporate bond market, who found a power-law distribution in client dealer links with an exponential tail. One-third of clients had only one dealer link, with a small fraction having a large number of links. Similarly, the UK gambling market can be categorized as an OTC market where bookmakers are the dealers and gamblers are the clients. Again, 44% of clients have only one account with a mean of 2.7 accounts, UK Gambling Commission (2019). These empirical observations of the network topology in the client-dealer network form the basis of our network modeling.

Although networks offer an excellent description of the topology of who is connected to who, the trading interactions or microstructure¹ determine the traded prices once the network has been set. While much of economics abstracts from the mechanics of trading, microstructure literature analyzes how specific trading mechanisms affect the price formation process, and how agents negotiate the terms of trade are critical. We assume a Bertrand price competition amongst the dealers with each client, and each dealer does not price discriminate due to not knowing the client's degree. This interaction could lead to the usual Bertrand paradox – when a client is connected to two or more market-makers, price competition will force the price to marginal cost causing zero profits in expectation. In financial market securities, we assume a common but unknown value to the security. This uncertainty in the common value causes competitive dealers to make non-zero bid-ask spreads and, importantly, non-zero pay-offs in equilibrium. This unknown common value effect on Bertrand competition is similar to the results of Spulber (1995), who analyzed the independent private values case and found a similar non-zero profit result. This non-zero profit in competitive equilibrium leads to the dealer's price to depend on the client base's degree of connectivity, which leads to equilibrium prices to depend on the network topology. This result is in contrast to many microstructure models, for instance, Kyle (1985), who postulated a zero-profit condition without any game-theoretic foundation but is in line with the models of Bernhardt and Hughson (1993) and Bilais et al. (2000), who also predicted positive profits in equilibrium.

Our network model consists of two elements: the nature of the price forming interactions between clients and dealers, which we assume is a standard Bertrand price competition with unknown common values. Secondly, the nature of the set of dealers that the client requests trading prices from, the network structure. Much of the literature uses search models for this stage, but our model uses a fixed network which better represents the nature of client-dealer links in many OTC markets, such as corporate bonds and derivatives. Empirically, while some markets can be described by random meetings among traders who are small relative to the market, in others,

¹Maureen O'Hara defines microstructure as the study of the process and outcomes of exchanging assets under explicit trading rules.

in particular, client-dealer OTC markets, relationships are not random. Due to relationship formation and maintenance costs, clients adopt relatively few links to dealers to find prices and trade. This formulation is in line with the assumptions of Babus and Kondor (2018) and others who rely on a fixed network and the empirical work of Hendershott et al. (2020) in the corporate bond market and Mallaburn et al. (2019) in the Bank of England empirical study on investment grade and high yield bonds.

The roadmap of this thesis is as follows: Chapter 1 provides an introduction and a review of the relevant OTC and microstructure literature, Chapter 2 introduces the microstructure model that governs the price formation process, Chapter 3 looks in detail at the nature of the equilibrium solutions in a duopoly dealer market, Chapter 4 introduces networks and specifies the network topology of the network model, Chapter 5 introduces the payoffs in a multi-dealer network game, numerical procedures for finding them and looks at the effects on the equilibrium in various informational regimes and network topologies, Chapter 6. presents an empirical study of the UK gambling market and applies the model to calculating bookmakers odds prices.

1.2 Related Literature

Our model can be split into two main parts. Firstly, the price negotiation process between individual customers and their connected dealers. Secondly, applying this microstructure to a network setting of multiple customers and dealers with heterogeneous connections. The first part is related to the traditional microstructure literature, and the second part is related to models of search on a network and specifically to the extensive work on trade in OTC markets.

1.2.1 OTC Literature

The standard method of analyzing OTC markets has been the search and bargaining approach, and the cornerstone to modeling decentralized markets was first proposed in the seminal paper by Burdett and Judd (1983). They proposed a model where customers search firm counterparties with an associated cost. They considered both nonsequential search, where consumers decided ex-ante how many firms to contact for

prices, and noisy sequential search, where customers continue searching until certain conditions are fulfilled. They show that an equilibrium strategy exists in both cases and leads to price dispersion in the market.

The most common approach to modeling price formation in OTC markets is based on this random search and matching approach, typically without a fixed network. This approach to modeling OTC markets was pioneered by [Duffie et al. \(2005\)](#); [Pedersen et al. \(2007\)](#), who study how asset prices in OTC markets are affected by the search costs associated with the search for counterparties and subsequent price bargaining. They show how bid-ask spreads are lower if traders can more easily find other traders or are connected to multiple market-makers and characterize the equilibrium pricing. They subsequently extend their search frictions analysis by looking at the effects of risk aversion and supply shocks.

Another important paper by [Lagos and Rocheteau \(2009\)](#) looks at the liquidity providing function of dealers and its importance during periods of market turmoil. Their model involves allowing continuous trading between dealers, but trading with investors is subject to search delays and bargaining. They examine the dealer incentives that satisfy both liquidity provisioning and market efficiency. In their later paper (2009), they extend the model of [Duffie et al. \(2005\)](#) by allowing unrestricted asset holding and show how asset demand changes are a key determinant of the bid-ask spread and other trading metrics. [Afonso and Lagos \(2015\)](#) develop a model of the federal funds market using the search and bargaining approach of Duffie, where banks have to search for a suitable counterparty and then negotiate the terms of the loan size and repayment schedules. [Duffie et al. \(2017\)](#) characterize the role of benchmarks in OTC markets as a mechanism of reducing informational asymmetry between dealers and customers, who, similar to our model, are typically less well informed than the dealers. They show that providing a benchmark in opaque OTC markets can improve efficiency by encouraging customer entry, improving matching efficiency, and search costs. This effect is also evident in our model by increasing the signal precision of the clients relative to the dealers.

By design, in OTC search models, transactions are between atomistic dealers through non-persistent links. Therefore, using a fixed network, our model's methodology better captures the effects of greater centrality of the few large dealers that

intermediate the bulk of the trading volume. Empirically, while some markets can be described by random meetings among traders who are small relative to the market, in others, in particular, client-dealer OTC markets, relationships are not random but fixed and persistent through time. Due to relationship formation and maintenance costs, clients adopt relatively few links to dealers to find prices and trade.

In our model, we take the market network structure (who trades with whom) as given and, in this sense, are closer to this separate strand of literature that views agents as interacting on a fixed network, such as the seminal work by [Kranton and Minehart \(2001\)](#). They looked at a network of buyers and sellers (in an IPV setting) and showed, using Hall's Marriage Theorem, the existence of a matching in a bipartite buyer-seller network. They use an ascending price auction representation of competition to solve for an equilibrium, which they show is efficient and pairwise stable. The paper demonstrates how buyers' and sellers' network structure is formed and includes the idea of link costs in the formation. Several of these ideas appear in our OTC model. Firstly, we restrict client networks to a fixed small number of links due to link formation and maintenance costs. Secondly, we also use a heterogeneous asset estimate for each agent; however, we use a common value for the payoffs, which makes the auction calculations more difficult. [Gale and Kariv \(2007\)](#) develop a financial network model with similarities to Kranton et al. In this paper, agents interact on a fixed network structure similar to our model, although they consider a unimodal network where agents are either buyers or sellers. [Blume et al. \(2009\)](#) model intermediation over a fixed network and finds that equilibrium strategies are network-dependent.

1.2.2 Microstructure Literature

The market microstructure literature is extensive (see [Biais et al. \(2005\)](#) for a full review) but is divided into primarily two distinct approaches - Inventory Models and Asymmetric Information Models. A key assumption in most microstructure models is the existence of a proportion of informed and uninformed traders. This proportion is generally held to be common knowledge. These trader classes can be viewed as a simplification of the client group's level of 'informedness', and this informedness

of the clients is covered in our model with the variance or precision of the trader signal error. Low (relative) signal error corresponds to low variance in the model and mimics the literature's informed traders. Conversely, high signal variance corresponds to less informed traders. Many of the models also assume a zero profit condition in competition which is an assumption we do not make. Indeed, the non-zero profit condition is critical to the network effects of the model.

Inventory Models

The inventory-based models assume that a market-maker adjusts the bid-ask spread in response to the asynchronous arrival of trades. These models assume that the market-makers' primary role is as liquidity providers and demonstrate how the bid-ask spread compensates them for price risk on inventory. [Garman \(1976\)](#) was the first to model microstructure from an inventory perspective. Briefly, the model considers a single market-maker and assumes that buyer's and seller's arrival follow a Poisson process. The arrival frequencies depend on prices traded, and the dealer uses a bid-ask spread to make profits causing lower amounts of trade with a wider spread. The dealer constraint is that inventories are targeted at a specific level. The asset is assumed to follow a mean zero random walk, hence ensuring a non-zero probability of bankruptcy over time. In order to counter this, market makers adjust prices dynamically with inventory positions.

Similar to Garman, [Amihud and Mendelson \(1980\)](#) extend the Garman model where the dealer's inventory is allowed to fluctuate between two bands. Quotes are updated when inventory approaches these bands to affect the arrival rate of buyers and sellers. Like Garman, traders are assumed to be uninformed liquidity traders, and dealers only include inventory in the pricing function.

[Ho and Stoll \(1981\)](#) focus on how risk-averse dealer's inventory, processing costs, and adverse selection determine the bid-ask spread. The additional assumptions in this model are; a dealer with risk aversion with no hedging possibilities and seek to maximize the final utility of wealth. This model concludes that the bid-ask spread is independent of the inventory position but is affected by the dealer's risk aversion.

[Ho and Stoll \(1983\)](#) extend the model and show that inventory affects the dealer's

quote level but not the spread magnitude is robust to multi-period and multi-dealer settings.

More complicated models have been introduced, and some of the major contributions have been: Stoll (1989), Huang and Stoll (1997), which add complications of adverse selection and, Bollen et al. (2004), who uniquely use option analysis to identify the components of the bid-ask spread. However, the setup continues with the assumption of informed and uninformed traders, where by definition, the informed traders are more informed than the dealer.

Asymmetric Information Models

Asymmetric information models assume that the market comprises agents with varying degrees of information regarding an asset's price. Again, the market-maker role is assumed to provide liquidity to traders who can be either informed or uninformed about the true asset value. Market-makers make losses when trading with informed traders and profits when trading with uninformed traders leading to the market-maker trying to recoup informed trader losses from uninformed traders by providing a bid-ask spread. Rational, competitive market-makers set their bid and ask prices accordingly, and more extreme information asymmetries lead to wider bid-ask spreads. These asymmetric information models are further classified as either strategic trade models, where clients participate in the market only once, or sequential trade models, where randomly selected clients sequentially arrive at the market.

Sequential trade models refer to the class of asymmetric information models where randomly selected traders arrive sequentially at the market. The main assumption in these models is the existence of heterogeneously informed traders. These are defined as "informed traders," who trade due to private information on the asset's fundamental value, and "liquidity traders," who trade for exogenous reasons, like portfolio rebalancing or liquidity needs.

Copeland and Galai (1983); was the first paper to consider the cost of asymmetric information. A simple setup of a proportion of informed and uninformed traders interacts with a dealer who makes a bid and an ask price. The ratios of the types of traders are known, and the dealers' optimization problem is simply the maximization of the expected payoffs from both types of traders.

Glosten and Milgrom (1985) produced a seminal model in the literature similar to the Copeland model. An asset has a high or low payoff with a certain probability, and the informed traders know the actual realization. The uninformed traders trade randomly. Notably, the dealer is uninformed but knows the proportion of informed traders and has a zero profit condition. As trades enter, the dealer updates the probabilities of the expected value of the asset. The model predicts that order flow is correlated, and bid-ask spreads decline over time as the actual probability precision narrows.

Easley and O'Hara (1987) extend the Glosten model, where both sets of traders can decide between large or small trade volumes. There are now two market-makers and a zero profit condition. The model's main point is that their bid-ask spreads are affected by trade size since large trades are correlated with adverse selection. In their follow-up paper in Easley and O'Hara (1993), they added a no-trade option for uninformed traders, thereby adding correlation to the timing of trades.

In Strategic trade models, traders and market-makers form expectations over the other's behavior in order to find an optimal strategy. The seminal model in this field is the model by Kyle (1985) and has been extended by various authors. In particular, Vives (2011), who formulates a model based on Kyle and uses dealer price schedules to examine equilibrium. Babus and Kondor (2018) then extend the Vives model to a network setting.

The Kyle (1985) model is set up similarly to the other asymmetric information models: a single-period model with a random (normally distributed) final value asset, a group of liquidity traders, a single risk-neutral insider trader, who knows the final realization, and a single risk-neutral uninformed market-maker. The liquidity traders submit orders to buy a random (normally distributed) amount of shares of the asset, and the insider must decide on optimal trade size. The market-maker must strategically decide on a bid-ask spread strategy that incorporates the probability of the insider trading. As order volumes increase, the market maker makes increasing spreads as the size of the aggregate order flow is correlated to the insider trades. The insider also models the market maker's strategy and submits orders accordingly. Admati and Pfleiderer (1988) provide an extension of the Kyle model by allowing some uninformed traders to behave semi-strategically by allowing them to time their trades.

They are referred to as discretionary traders because although they are obliged to trade, they have some discretion on their trades' timing. Our model's clients bear some similarity to these discretionary traders, but instead of timing, they have a reservation price, based on the true value, that must be met before they can trade.

Chapter 2

A Model of Trading in OTC Markets

An OTC market is distinguished from a centralized market by the network of connections of the participants - who is connected to whom. As clients can only receive trading prices from connected dealers, the form of the interactions or trading protocols between the clients and dealers and their link topology can lead to different prices trading in different parts of the network. We model the interaction between clients and dealers as a price competition process in the style of [Bertrand \(1883\)](#) with unknown common values between the clients and their connected dealers. We begin by defining the agents, properties of the asset, their connectivity in a network, and finally, the trading protocols that govern the price-setting process.

2.1 Agent Types

There are two agent types in our OTC market model - clients and dealers, which are analogous to the firms and customers in traditional economic theory.

2.1.1 Clients

Clients have an exogenous reason for trading (and can be either a buyer or a seller). They are boundedly rational in the sense that they attempt to estimate the true value but naively use this estimate as their reservation price before accepting a dealer quote. This client specification is in line with [Gode and Sunder \(1993\)](#), who found

that the artificial traders that were the most similar to human traders were the zero-intelligence ones with budget constraints, i.e., traders who were not allowed to trade at a perceived loss. In our model, each client's reservation price is a random variable that depends on the true value and sets the boundedly rational strategy of the clients.

This formulation of the client reservation prices allows us to model differences in information sets regarding the true value via signal variance differences and is similar to the fundamentalist client specification in [de Kamps et al. \(2012\)](#), who also used client valuations to determine their trading strategy.

2.1.2 Dealers

The dealers are risk-neutral strategic agents who seek to maximize their expected profits by quoting a buy price P^B and a sell price P^A to the clients. They are strategic in the sense that they attempt to estimate the true value of the asset traded, form beliefs as to the competitor estimates, and use these estimates to construct their buy and sell prices. Dealers fulfill the liquidity-providing market function by being compelled to both make a price if requested and trade on that quoted price. These assumptions contrast with the pure uninformed liquidity provider assumptions of both the inventory and Asymmetric information-based microstructure models. They set their buy prices and sell prices to maximize their expected profits in equilibrium, and the difference between their buy and sell prices is known as the bid-ask spread. Dealers base their bid and ask prices as a function of their estimate of the true value and an expectation of all other agents' estimates and assume that other dealers will be doing the same.

2.2 Common Asset True Value and Signals

In many OTC markets, such as derivatives and corporate bonds, the trading frequency in each issue and the proportions of buyers and sellers at any one time make matching trades between clients difficult, suggesting dealers need to have a valuation of the final payoff of the asset. For example, in the US and UK corporate bond market, each individual bond issue trades on average 2.4 times per day, [Bessembinder et al. \(2020\)](#). For illiquid issues, they trade even less often. In the corporate bond market in the EU,

there are over 150,000 outstanding securities issued where only the top 3000 trade more than once per day. In addition, average trade sizes are over \$1million compared to equities of \$43,000, International Capital Markets Association. Dealers knowingly assume the opposite side of the client trade and receive any ensuing payoffs from the asset. Therefore, they have an incentive to value the asset as accurately as possible. In our model, we assume a common unknown value of the asset to each agent at the time of trading. For most traded financial assets, the final or true value is the same for all agents. However, the process of measuring this value a priori is subject to both measurement and model error.

The clients and dealers estimate the asset's value before trading but with some error or uncertainty, which we characterize with some noise $N(0, \sigma)$. Agents attempt to measure this true cost (value), and we can model this by each of them receiving a sample (signal) from a known distribution. This measurement of the true value is subject to independent and idiosyncratic measurement error or interpretation, and we use a normal distribution as the common distribution of the error signals. In many fields, measurement errors (as opposed to systematic errors) are modeled as normally distributed. We can view measurement as the result of a process, where each step in the process may lead to a small error with a probability distribution (F.P. Schloerb, Computational Physics, University of Mass). The sum of these errors over all of the measurement steps leads to a final error that is normally distributed, whatever the error distribution in the individual steps.

In our case, since the asset's true value is subject to many micro-measurements that may come from any distribution, we can apply a normal distribution. For instance, the price forecast for a stock is dependent on correctly measuring current earnings, likely sales growth, consumer trends for their products, global interest rates, and general economic conditions, to name but a few. The central limit theorem says that the distribution of these price estimates will tend to a normal distribution as the number of estimated factors grows large.

There is some empirical evidence for investor valuations being normally distributed: CXO Advisory group analyzed 6582 forecasts¹ for the US stock market between 2005 - 2012, and found that their forecast accuracy was on average close to

¹see <https://www.cxoadvisory.com/gurus/> for the full analysis

zero with an approximately normal distribution. Also, the estimates of Vote shares in forthcoming elections are modeled with a normal distribution by polling companies in Chapter 6.

In line with Babus and Kondor (2018) and others, we have adopted an additive (normal) error term as the signal and the statistical form of the beliefs of the true value. Although stock price and earnings forecasts are not an exact proxy for beliefs (due to herding and other strategic behavior), SPX forecasts were found to be approximately normally distributed with a zero mean (true value), and Kim et al. (2017) analyzed individual stock analyst forecasts using IBIS data over a 30 year period. They found a lognormal distribution best describes the forecasts with a positively biased mean and unbiased median. Individual forecasts are complex as they are affected by the individual incentives of the financial analysts, as evidenced by some fat tails in extreme forecasts. In any event, the normal distribution is a good approximation to a lognormal distribution when the variance is small relative to the mean value. A widely used heuristic rule is that if X is $LN(m,s)$, then X is approximately $N(m,s)$ if $m > 6s$. Suppose a security has a mean value of 100 with an estimation error or 10 or 15%, then figure 2.1 illustrates the approximation.




FIGURE 2.1: Illustration of the normal approximation to lognormal when means are greater than standard deviations

As an example of financial market instrument measurement error even in the simplest of instruments, Cammack (1991) found that empirical bond auction results suggest that imperfect information is present in the Treasury bill market. This financial instrument should have one of the lowest forms of potential measurement

or interpretation error (short duration and limited pricing factors). She found the mean auction price for 3-month bills was, on average, four basis points below the comparable secondary market price for the 1973-84 period. She concluded that this "downward biasing" is positively related to the anticipated amount of dispersion of auction bids that suggested that auction bidders use a bidding strategy that accounts for their lack of agreement about the bill's value.

Suppose the true asset value is V so that the client i has an estimate V_i which we model by them receiving a signal $V_i = V + \epsilon_i$, where $\epsilon_i \sim N(0, \sigma_T)$. Similarly for the signal of dealer j : $V_j = V + \epsilon_j$ where $\epsilon_j \sim N(0, \sigma_M)$. Note that $E[V_i] = V$, the true value, for both clients and dealers. We assume that all dealers have the same signal error variance, although relaxing this assumption does not fundamentally change the model or the results. The variance of the signal could also be interpreted as the sophistication of the clients. If their level of sophistication is anti-correlated with their connectivity, this may add an important multiplying effect. We only consider that clients all have the same signal variance in the model and analyze the dealer behavior towards them. However, if client degree and signal variance are anti-correlated, certain effects will be magnified, and these will be clear in the later sections.

2.3 Client and Dealer Strategies

As above, dealers and clients begin the trading process by attempting to estimate the true asset value. The clients set their reservation prices based on their signal level and have a naive strategy that maps each possible signal (value estimate) onto a reservation price. The dealers make bid and ask prices around their own signal. We assume a symmetric linear bidding strategy from the dealers, which is justified by the equilibrium solutions of [Wilson \(1969\)](#) in the 2-bidder case and both [Levin and Smith \(1991\)](#) and [Wilson \(1992\)](#) in the n -bidder case and explored in section [2.9](#) and [3.3.2](#). That is, each dealer i receives a signal V_i regarding the true value V , $V_i \sim N(V, \sigma_M)$ and makes a price quote of $(V_i - \delta_i, V_i + \delta_i)$ to their connected clients, where δ_i is known as the semi bid-ask spread. This linear bidding strategy maps each

signal (asset valuation) onto a trading price.

$$\begin{aligned} P^{BID}(V_i) &= V_i - \delta_i : V_i \rightarrow \mathbb{R} \\ P^{ASK}(V_i) &= V_i + \delta_i : V_i \rightarrow \mathbb{R} \end{aligned} \tag{2.1}$$

Where, $V_i \in (-\infty, \infty)$, $\delta_i \in [0, \infty)$.

We assume that the dealers are strategic in their selection of δ_i . Their buy and sell prices incorporate a reduction δ_i to their estimate V_i that is designed to compensate for the possible error in their estimate, a profit margin, and an expectation of other dealer strategies. This amount over the estimated true values δ is the semi bid-ask spread. The dealers do not see the other dealers' prices when they make their bid and ask prices. The technical justifications for a linear bidding strategy are discussed in the next chapter.

2.4 Connectivity in a OTC Financial Market

In OTC markets, empirical observations have shown the network structure to have three main features. Firstly, a highly connected core of dealers, [Li and Schürhoff \(2019\)](#) and others. Secondly, a small subset of dealers that intermediate a disproportionate amount of trade, [Hendershott et al. \(2020\)](#), [Mallaburn et al. \(2019\)](#) and others, and thirdly, client nodes with relatively few but persistent links to the dealers, [Babus and Kondor \(2018\)](#). These empirical observations of the OTC network topology (detailed in section [4.3](#)) suggest a core-periphery network structure. However, we do not explicitly model the dealer-dealer subnetwork and focus only on the client-dealer subnetwork.

The competitive process in OTC client-dealer markets is comprised of two economic agent types - clients and dealers. The clients connect to the dealers in some way and request quotes from these connected dealers. Typically, clients have only a few links to dealers, as demonstrated by the empirical work of [Hendershott et al. \(2020\)](#) in the corporate bond market. Clients do not typically connect with themselves due to search costs constraints. In OTC financial markets, dealers often also connect to other dealers in order to trade together due to heterogeneous inventory

capacities, Chung-Yi Tse et al. (2021). The dealer-dealer behavior is not modeled, but we assume that dealers are connected in a dense network, Mallaburn et al. (2019), Li and Schürhoff (2019) (and many others) and use this inter-dealer network to trade with each other in order to balance inventory, meet liquidity needs and speculate, de Kamps et al. (2012).

We construct a simplified version of an OTC client-dealer market by representing it as a bipartite network (a network with two node types), where nodes represent the dealers and clients and the links represent a possible trading relationship between them. Suppose there are N dealer nodes and M client nodes.




FIGURE 2.2: A graphical bipartite network representation of an OTC market

Figure 2.2 shows an example of a bipartite network representation. Dealers D_i have a valuation V_{D_i} of a homogenous asset and make prices P_i to their connected clients C_j . The clients have a valuation V_{C_j} of the asset. Each client can only trade with the dealer(s) that they are connected to at a price P_i suggested by that dealer. The client selects the best price of their connected dealer(s), and trade occurs. The exact trading protocols are discussed in detail in section 2.6.

2.5 OTC Trading Architecture

Most OTC dealer markets have some standard trading protocols, and Biais et al. (2006), after extensive interviews in the OTC client-dealer market for bonds, identified several essential features of OTC dealer markets, which guide the trading protocols of our model:

- Clients request quotes from dealers.

- Dealers respond to these requests for quotes by posting bid or ask prices simultaneously and independently. These quotes are firm, and the customer allocates the order to the best quote.

- There is no pre-trade transparency, in the sense that the customer does not see quotes before submitting the request for quotes. Neither do the dealers see the quotes of their competitors.

This schedule of trading interactions can be modeled as a simple Bertrand price competition with an unknown common 'true' value V . The dealer with the best price, P , trades with the client and receives a payoff of $P - V$. All other dealers who are also connected to this client receive a zero payoff. We assume that the dealers are simultaneously pricing with other clients and the dealers do not price discriminate and make the same prices to all their connected clients.

The model of the OTC market that we use can be described as a 1-period, quote driven market with a single risky asset traded that has an unknown common value a priori to all agents. Price sensitive, price-taking clients interact through risk-neutral price-setting dealers who are strategic in their price setting and who seek to maximize their payoffs in equilibrium. The dealers make executable 2-way prices (buys and sells), and the clients contact the dealers and accept their quotes to trade or not, depending on their reservation prices. We assume a unit demand from the clients for the asset.

2.6 Model Protocols

- We assume that it is common knowledge that each dealer and client estimate the true value of the asset with a known (normal) probability distribution. Furthermore, it is common knowledge that each dealer adds an linear amount to their estimate to make a trading price

- Dealers estimate true asset price V , modelled by each receiving a signal $V_i \sim N(V, \sigma_M)$

- Clients C_j estimate true asset price V , modelled by receiving a signal $V_j \sim N(V, \sigma_T)$

- Dealers construct bid and ask prices based on their estimate of $P_i^B = V_i - \delta_i$

and $P_i^A = V_i + \delta_i$

- Clients form reservation prices R based on their signal of $R = V_j$
- Clients C_j request trading prices from their connected dealers $N[C_j]$
- If best trading price from connected dealers is better than R , then a unit trade occurs at the best dealer's price
- After all trading, true value V is realized and payoffs are calculated

A client asking N dealers for a price in an OTC market, where the dealers do not see each other's quotes, is strategically equivalent to the client conducting a blind first price common value sealed bid auction with the N dealers. In this setup, the client is the auctioneer, and the dealers the bidders.

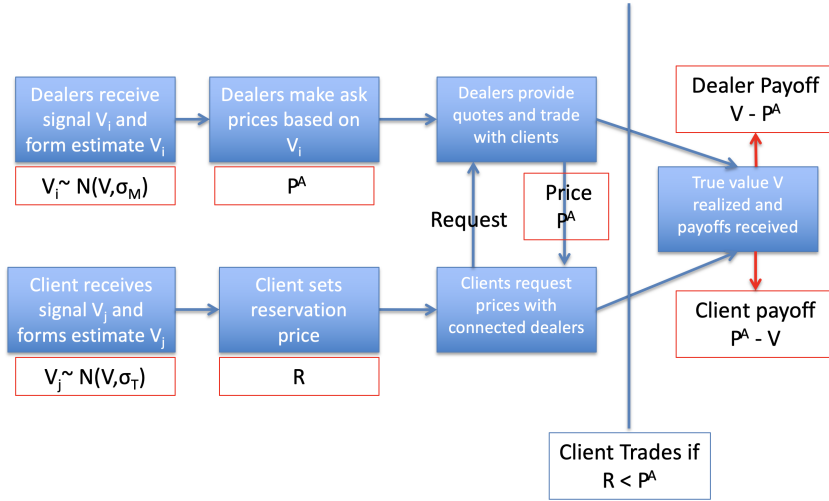


FIGURE 2.3: The trading protocols for a buying client connected to a subset of dealers

2.7 Payoffs

We begin by examining the base case of the expected payoffs for multiple competitive dealers and one client. This is then extended to the M client case, where the clients are connected to the dealers in a sparse network structure which we define in detail in Chapter 4.

Let there be N dealers and 1 client. Dealers and clients receive independent signals $\{V_i\}$ as to the true value of the asset with a signal error that is normally distributed with mean V and standard deviation σ_M for each of the dealers and σ_T for the clients.

Clients set their reservation price at their signal - a naive estimation, to model the boundedly rational behaviour of clients. We make the standard assumption that dealers seek to maximize their expected payoffs in equilibrium.

2.7.1 N dealers, 1 Client payoffs

In an OTC market with one client and n dealers, the payoff for dealer i, in competition with the N-1 other dealers from a single selling client is :

$$\Pi_i = \left\{ \begin{array}{l} V - P_i : \text{if } P_i > P_j \forall_j \text{ and } P_i > R \\ 0 : \text{otherwise} \end{array} \right\} \quad (2.2)$$

where P_i is the buying quote of dealer i and V is the true value. P_j represents the buying quotes of the other N-1 dealers. R is the client reservation price that must be bettered if a trade is to occur and is a random variable dependent on the client's signal variance. The dealer has a non-zero payoff only if they have the highest price with an analogous expression for the client buying. The dealer with the highest price gets selected to trade, and the expected payoff is the expected value of this highest price, conditional on the price being a maximum for buys (and minimum for sells).

Client payoffs are the opposite of the dealer payoffs, and the game is zero-sum; hence the client payoff is:

$$\Pi_{client} = \{P_{\max} - V : \text{if } P_{\max} = \text{Max}(P_i) > R : \text{zero otherwise} \} \quad (2.3)$$

There is an analogous expression for a buying client. In this case the dealers make a buying price of $P_i = V_i - \delta_i = V + \epsilon_i - \delta_i$, that is, they reduce their signal by δ_i and the payoff for dealer i becomes:

$$\Pi_i = \{-\epsilon_i + \delta_i : \text{if } \epsilon_i - \delta_i > \epsilon_j - \delta_j, \forall_j \text{ \& } P_i > R : \text{zero otherwise} \} \quad (2.4)$$

The dealer with the highest signal error, adjusted by their bid-ask spread, gets selected to trade. This bidding strategy can lead to a negative payoff problem if δ_i is too low and leads to the widely observed phenomena of the winners curse².

²Common Value Auctions and the Winners' Curse, Princeton U Press (2001)

This phenomenon occurs in common value auctions when bidders do not sufficiently compensate for the fact that the probability of having an incorrect price is less than the probability of an incorrect price conditional on being the best price. This non-strategic behavior can lead to overbidding and negative payoffs.

The expected payoff for buys is the expected value of $-(\epsilon_i - \delta_i)$ conditional on this being a maximum in the game and for sells is the expected value of $-(\epsilon_i + \delta_i)$ conditional on this being a minimum in the game. When ϵ is from a symmetric normal distribution with a mean zero, these payoffs are identical. The expected value is calculated below for the n-player case.

We assume each dealer D_i , uses a linear bidding function (strategy) that maps their signals S_i (estimates) onto a trading price and quotes a selling price (after receiving their signal) of $P_i = S_i + \delta_i = V + \epsilon_i + \delta_i$ and has a payoff of $(\epsilon_i + \delta_i)$ if $(\epsilon_i + \delta_i)$ is less than all the other $(\epsilon_j + \delta_j)$ of the other dealers and also less than the client reservation price $\epsilon_T + R$, zero otherwise. Each $(\epsilon_i + \delta_i)$ is normally distributed as $N(\delta_i, \sigma_i)$.

Consider the N-1 other dealers and the client reservation price. We first need to calculate the probability that dealer i's price is less than the minimum of this group.

Formally we are looking for the probability of dealer i, to win a competitive auction with N-1 other dealers subject to a client reservation price. $P(p_i = V + \epsilon_i - \delta_i < Y)$ where $Y = \min\{\epsilon_1 + \delta_1, \epsilon_2 + \delta_2, \dots, \epsilon_{i-1} + \delta_{i-1}, \epsilon_{i+1} + \delta_{i+1}, \dots, \epsilon_N + \delta_N, \epsilon_T + R\}$.

Consider N+1 independent normal random variables, X_0, X_1, \dots, X_N , where each X_i has a mean δ_i and standard deviation σ_i .

Hill, J (2011), analysed the case of the probability of an independent normal RV $X_0 \sim N(\delta_0, \sigma_0)$ being less than the minimum Y of N other normal (independent) RV's X_1, \dots, X_N with $X_i \sim N(\delta_i, \sigma_i)$ is :

$$Pr(X_0 < Y) = \int_{-\infty}^{\infty} \frac{1}{\sigma_0} \phi\left(\frac{s - \delta_0}{\sigma_0}\right) \prod_{j=1}^N (1 - \Phi\left(\frac{s - \delta_j}{\sigma_j}\right)) ds \quad (2.5)$$

where δ_j is the spread charged by dealer j who receives a signal $S_j \sim N(V, \sigma_j)$ and $\Phi(\cdot)$ and $\phi(\cdot)$ are the distribution and density functions of a standard normal distribution.

The expected value of this RV is:

$$E[X_0] = \int_{-\infty}^{\infty} s \Pr(X_0 = s) \Pr(Y > s) ds \quad (2.6)$$

Therefore, we can express the expectation of the payoff Π_i from each buying client as:

$$E[\Pi_i] = \int_{-\infty}^{\infty} \frac{s}{\sigma_{Mi}} \phi\left(\frac{s - \delta_i}{\sigma_{Mi}}\right) \left[\prod_{j=1, j \neq i}^N \left(1 - \Phi\left(\frac{s - \delta_j}{\sigma_{Mj}}\right)\right) \right] \left(1 - \Phi\left(\frac{s + R}{\sigma_T}\right)\right) ds \quad (2.7)$$

Now consider the selling clients. The calculations are the same but now involve maximum order statistics. Specifically, each dealer makes a buying price (after receiving their signal) of $V + \varepsilon_i - \delta_i$ and has a payoff of $P - V = -(\varepsilon_i - \delta_i)$ if $(\varepsilon_i - \delta_i)$ is greater than all the other $(\varepsilon_j - \delta_j)$ of the other dealers and greater than the client reservation price $\varepsilon_T + R$, zero otherwise.

Consider the $N-1$ other dealers and the client reservation price. We are interested in the probability that dealer i 's price is greater than the maximum of this group. By the same procedure as above, this expected payoff equals:

$$E[\Pi_i] = - \int_{-\infty}^{\infty} \frac{s}{\sigma_{Mi}} \phi\left(\frac{s + \delta_i}{\sigma_{Mi}}\right) \left[\prod_{j=1, j \neq i}^N \Phi\left(\frac{s + \delta_j}{\sigma_{Mj}}\right) \right] \Phi\left(\frac{s - R}{\sigma_T}\right) ds \quad (2.8)$$

It is easy to show that the payoff functions (2.7) and (2.8) are the same. This is also apparent by a symmetry argument. In this chapter we are only considering the case of $\sigma_{Mi} = \sigma_M$, that is, all dealers have same signal variance.

Variance of the expected payoffs from a single client can also be calculated as:

$$V[\varepsilon_i - \delta_i] = E[(\varepsilon_i - \delta_i)^2] - E[(\varepsilon_i - \delta_i)]^2$$

$$V[\varepsilon_i - \delta_i] = \int_{-\infty}^{\infty} \frac{s^2}{\sigma_{Mi}} \phi\left(\frac{s + \delta_i}{\sigma_{Mi}}\right) \prod_{j=1}^N \Phi\left(\frac{s + \delta_j}{\sigma_{Mj}}\right) \Phi\left(\frac{s - R}{\sigma_T}\right) ds - \left(\int_{-\infty}^{\infty} \frac{s}{\sigma_{Mi}} \phi\left(\frac{s + \delta_i}{\sigma_{Mi}}\right) \prod_{j=1}^N \Phi\left(\frac{s + \delta_j}{\sigma_{Mj}}\right) \Phi\left(\frac{s - R}{\sigma_T}\right) ds \right)^2 \quad (2.9)$$

In summary, with one client (buyer or seller) that connects to a sub-network of dealers D_1, \dots, D_k , and sets a reservation price of $V_j \sim N(V, \sigma_T)$ and each dealer uses a symmetric bidding strategy of $(V_i - \delta_i, V_i + \delta_i)$, $V_i \sim N(V, \sigma_M)$ the expected payoff to dealer i is:

$$E[\Pi_i] = - \int_{-\infty}^{\infty} \frac{s}{\sigma_M} \phi\left(\frac{s + \delta_i}{\sigma_M}\right) \left[\prod_{j=1, j \neq i}^k \Phi\left(\frac{s + \delta_j}{\sigma_M}\right) \right] \Phi\left(\frac{s}{\sigma_T}\right) ds \quad (2.10)$$

This is the expected payoff to each dealer and depends only on the choice variable δ_i , the semi bid-ask spread, and the variance of the client and dealer estimate σ_M, σ_T .




FIGURE 2.4: Comparison of numerical simulation to the derived payoff functions

As a check, we ran numerical simulations using vectors of length z of normally distributed random variables. We compared the payoff to each dealer in each row for various different spread and variance levels. As an illustration, we look at 3 dealers and 1 client. All dealers receive a signal $N(V, 2)$ and the client a signal $N(V, 3)$. Dealer 1 uses a spread of 2, dealer 2, 3 and dealer 3, 4. For each random signal realization we select the dealer with the highest signal and set the payoff to the actual payoff or zero if their price was not the highest. The average payoff for each dealer over each sample set was taken and compared to the theoretical level above. The results

are illustrated in 2.4. Figure 2.5 shows a sample of the 5,000,000 auctions simulated for the average values with the winner highlighted in red. A client win means the reservation price was not bettered and no trade occurred.

Signals - Spread				
	dealer 1	dealer 2	dealer 3	client
auction 1	-2.2445219	-4.9191259	-0.8965697	-5.3693106
auction 2	-2.1210979	-5.2290228	-2.6749462	-0.3996183
auction 3	-6.7663152	-4.2198658	-3.8345943	2.4476034
auction 4	-2.0090581	-5.5152121	-3.1923464	2.1910921
auction 5	-3.960611	-1.1157963	-5.4758354	2.4743563
auction 6	-0.4500487	-6.3949648	-0.4949117	-1.5237575
auction 7	-5.1752558	-2.7101919	-5.3368575	4.4597076
auction 8	-2.9549785	-2.8102067	-2.0885101	2.0532034
auction 9	3.1310554	-1.3946846	-1.0704334	1.1527375
auction 10	-4.4629785	-3.6133331	-3.1463319	-0.6003823
auction 11	-2.6168425	-0.402491	-5.1034555	-3.0523871
auction 12	-0.8422609	-3.262509	-5.2978388	0.1814647

FIGURE 2.5: Sample of random realizations in each auction

2.8 Bid-Ask Spreads, Bertrand Competition and Auctions

2.8.1 Bid-Ask Spreads and Bertrand Competition

In the traditional Bertrand model of competition, Bertrand (1883), n firms, each with a marginal cost C , set selling prices P_i of a homogenous good to a client who selects the lowest price to trade. This firm with the lowest price captures the whole market and makes a payoff equal to $P_i - C$. This leads to a Nash equilibrium solution of all firms pricing at marginal cost, $P_i = C$, in competition.

In an OTC market, the dynamics are the same, but in addition, dealers simultaneously provide a buying price for the asset. Suppose there are n dealers connected to a client. Dealers ($i = 1 : n$), make selling prices, P_i^{sell} , to a client who buys at the lowest price, $\min(P_i^{sell})$ of their connected dealers. Simultaneously, dealers make a buying price P_i^{buy} to a client, who selects the highest price $\max(P_i^{buy})$ of their connected dealers. The payoff to the winning selling dealer is $P_i^{sell} - V$, where V is the

true or realized value of the asset. Similarly, for the buying dealer, they have a payoff of $V - P_i^{buy}$.

If the true value V is known, then in competition, the standard Bertrand results imply that each dealer would price both buy and sell prices at the marginal cost V , causing a zero bid-ask spread. However, this common value V is not known ex-ante to making a price. We need the Nash equilibrium solutions of a Bertrand competition with unknown common values in order to generate the bid-ask spread. When the cost is unknown, bidding your signal leads to the winner's curse problem - the dealer with the worst estimate gets selected to trade and has a negative expected payoff. Theoretically, this estimate needs to be 'shaved' by some amount to compensate for this.

2.8.2 Bertrand Competition and Auctions

A customer requesting a buying price from n firms, where the customer sells to the highest price, is the same as a price auction, where the customer is the auctioneer and the firms the bidders. In the OTC market case, there is a common value asset, firms do not see the quotes of other firms, and the highest price firm wins the trade. This is strategically equivalent to the client conducting a first price common value sealed bid auction with the n dealers. In this setup, the client is the auctioneer, and the dealers the bidders. Because of the unknown common value, the equilibrium results of marginal cost pricing in the standard Bertrand model do not hold, however this auction representation of the competition process can be analyzed using the equilibrium results of [Wilson \(1969, 1977\)](#) for the pure common value case and [Milgrom and Weber \(1982\)](#) for the more general payoff case. These equilibrium results of bidding strategies will equate to a bid-ask spread around the true value of an asset. This bidding strategy 'shaves' an optimal amount from the estimate in order to generate a bid and an ask price.

The next chapter looks in detail at a two dealer duopoly OTC market and looks at the nature of the equilibrium solution in a Bertrand unknown common value competition using the results of the [Wilson \(1969\)](#) equilibrium bidding strategies in common value auctions. The duopoly problem adds more insight than the general case into

the nature of the equilibrium solution as most of the equations have neat analytical solutions.

2.9 Equilibrium Bidding Strategies in First Price Common Value Auctions

2.9.1 Introduction

Several theoretical and empirical studies have focused on bidding behavior in the symmetric first-price common-value sealed-bid auction, FPSBA, given different empirical specifications of the distribution of bidders' signals and the distribution of the true value (e.g. [Thiel \(1988\)](#); [Levin and Smith \(1991\)](#); [Paarsch \(1992\)](#); [Wilson \(1992\)](#)). This appendix follows the assumptions and methodologies of [Levin and Smith \(1991\)](#) in deriving the linear form of the equilibrium pricing strategy in an FPSBA when the seller adopts a naive reservation price strategy. We show that the complication of a naive client reservation price strategy in an n -bidder FPSBA, where the client sets the RP at their estimate, is equivalent to an $(n+1)$ -bidder auction where the $(n+1)^{th}$ bidder uses a naive bidding strategy. This asymmetric variation has a similar solution as the symmetric variation described by Levin, Wilson, and Milgrom.

2.9.2 Setup

We assume that n risk-neutral bidders are bidding for a particular object, where the value of the object V is identical but unknown to all bidders prior to bidding. Before the auction, each bidder receives a private signal X_i concerning the value of the object. The seller sets an (unannounced) reservation price set at their own estimate X_c , and the bidder who submits the highest bid that is also greater than the reservation price is selected as the winner and awarded the object. Each bidder uses their private signal X_i to form an estimate of V , i.e., $E(V|X_i)$. No bidder knows the estimate of any other bidder. The bidders' signals X_i are positively correlated (affiliated) with a cumulative distribution function $F(X|V)$. The bidders have prior beliefs about the true value V which is characterized by the cumulative distribution function $G(V)$.

The number of bidders and the distribution functions of X_i and V are assumed to be common knowledge.

2.9.3 Solution to the Standard FPSB Auction

Levin and Smith (1991), in a comment on Thiel (1988) derive the explicit form of the equilibrium bidding solution in a standard common-value auction (no reserve prices). Each bidder observes a random signal X_i with variance σ^2 , which is an estimate of the unknown value V of the item being auctioned. $F(X_i|V)$ represents the conditional distribution of the i^{th} bidder's estimate, and $g(V|X_i)$ represents the bidder's posterior density for V given the signal X_i . Let $B(X)$ represent the symmetric equilibrium strategy used by each of n bidders, with $B'(\cdot) > 0$ such that the inverse $r(B) = B^{-1}$ is well defined. If all other bidders are using this strategy, then the i^{th} bidder's problem may be written as:

$$\max_{B(X)} \int_{\Sigma V|X} (V - B) F(r(B)|V)^{n-1} g(V|X) dV \quad (2.11)$$

Taking the first-order condition and additionally imposing the following three conditions proposed by Thiel (1988):

- 1) Each bidder's prior distribution of value V is diffuse: $g(V)$ is constant for all V .
- 2) Estimation errors are statistically independent of the item's true value: $F'(X_i - V|V) = f(X_i - V|V) = f(X_i - V)$
- 3) Each bidder's estimate of the value is unbiased: $E(X_i) = V$.

Leads to the differential equation:

$$B'(X) + K_1 B(X) + K_2 - K_1 X - 1 = 0 \quad (2.12)$$

This has a solution of:

$$B(X) = X - \frac{K_2}{K_1} + \beta \exp(-K_1 X) \quad (2.13)$$

Given the assumptions above and, additionally, the assumption that the estimation errors are normal and independent of the true value, then Levin and Smith

(1991), show that the optimal bid function in the first-price auction can be derived as a family of solution with a parameter $\beta \geq 0$

$$B(X) = X - \alpha_n \sigma + \beta \exp\left(-\int_{-\infty}^{\infty} t dG(t)^n X / \sigma\right) \quad (2.14)$$

Where $\alpha_n = \int_{-\infty}^{\infty} t^2 dN(t)^n / \int_{-\infty}^{\infty} t dN(t)^n$ is the ratio of the second and first moments of the maximum values of a standardized normal distribution. $\beta = 0$ corresponds to a linear solution.

Wilson (1992) remarked that in practice, bidding strategies are often constructed on the assumption that the marginal distribution of the common value has a large variance. For example, suppose that each estimate X_i has a normal conditional distribution with mean V and variation σ^2 , and that the marginal distribution of V has a normal distribution with variance s_0^2 . Then with no reservation prices, the limit of the symmetric equilibrium bidding strategy as $s_0 \rightarrow \infty$ is $B(X) = X - \sigma a_n$. That is, bidding strategies converge to linearity as normally distributed priors become diffuse in the limit.

Empirically, a linear bidding function in FPSBA has been observed in the US Treasury market, Cammack (1991) and Paarsch (1992) found that a value of $\beta = 0$ (a linear bidding function) was statistically the most likely value of β from the data of 144 separate auctions.

For both theoretical and empirical reasons, we focus on the linear bidding strategy solution in a standard FPSBA³.

2.9.4 Equivalence of Reservation Prices and Extra Bidder

Although in an FPSB auction with uninformative prior distributions and normal estimates, the symmetric Nash equilibrium strategy is linear in signals; this may not necessarily hold when we add the complication that the client also receives a private signal X_c about the true value V and adopts a naive reservation price strategy $RP(X_c) = X_c$. In this case the seller's ask price can be regarded as another bid, since

³From Levin and Smith (1991) "Linear strategies are not unknown in the literature. Previous researchers have noted their existence under the restrictions of 1-3 above. Richard Engelbrecht-Wiggans and Robert J. Weber (1979) were among the first to discuss the necessity of a diffuse prior (restriction 1). Michael H. Rothkopf (1980) and Robert L. Winkler and Daniel G. Brooks (1980) provided examples in which independent estimation errors (restriction 2) were used to derive linear bidding functions."

bidder i wins the standard n -player auction if $\{B_i > B_j \forall j\}$, $j = 1 \dots n$ and wins in a reservation price auction if $\{B_i > B_j \text{ and } B_i > B_c\}$ which is equivalent to an $(n+1)$ bidder standard auction with no reservation prices, where the $(n+1)^{th}$ 'bidder' uses a naive bidding strategy $B(X_c) = X_c$.

Although there are very few general results for asymmetric information cases, the asymmetry in this auction is solvable since the $(n+1)^{th}$ bidder is not strategic.

2.9.5 Signal and Bid Distributions

In the $(n+1)$ player auction, each bidder $i = 1, \dots, n$ receives a sample X_i from a distribution described by a CDF $F(V, \sigma_M^2)$ with mean V and variance σ^2 and the $(n+1)^{th}$ bidder (the seller) receives a signal from the same distribution F , with mean V but with variance σ_T^2 . Each bidder uses a bid function $B_i(X_i)$ which we assume is monotonically increasing in X .

Suppose bidder i wins the auction and we want to calculate the probability that bidder i was the highest bid. The set of bids $\{B_1(X_1), B_2(X_2), \dots, B_n(X_n), B_i(X_i)\}$, without knowing $B(\cdot)$, makes impossible forming the distribution of the maximum $B_i(X_i)$. Note in particular that $X_i > X_{n+1}$ does not imply that $B_i(X_i) > X_{n+1}$.

Bidder i wins the auction only if his bid was the largest out of the other bidders' bids and also larger than the $(n+1)^{th}$ bidders bid (which by design is their signal). We can assume a symmetric bidding function $B(\cdot)$ for the n bidders (by homogeneity), and since the bidding function is monotonically increasing in the signal, then we can work backward to the order of equivalent signals of the $(n+1)$ bidders by applying an inverse function $B^{-1}(\cdot)$ to all the bids including the seller. The probability of bidder i winning the auction is now the same as the probability of bidder i having the largest signal in this associated signal set.

The $n+1$ signal set now has the same ordering as the bid set and implies that if $X_i > X_{n+1} \implies B(X_i) > X_{n+1}$. Since $B(\cdot)$ is an unknown function, we still cannot calculate the maximum order statistics needed to derive the probability of winning; however, we know the exact distribution, F of these associated signals (bar one), and so bidder i wins the $(n+1)$ FPSB auction when his signal is largest in the associated signal set $S = \{X_1, \dots, B^{-1}(X_{n+1}), X_i\}$. The $\{X_j : j = 1, \dots, n\}$ are all distributed as F with mean zero and variance σ_M^2 , but $B(\cdot)$ is still an unknown function whose

inverse is assumed to exist. Therefore $z = B^{-1}(X)$ may not be described by $dist^n F$ (depending on the bidding function B of the n bidders and the nature of F). When F is normal, and B is linear, a particularly simple form emerges.

2.9.6 Equilibrium Bidding Strategy

Following the derivation of [Levin and Smith \(1991\)](#), we derive the equilibrium bidding strategy of the $(n-1)$ strategic bidders in a standard n -bidder FPSBA when bidder n uses a naive bidding strategy of $B(X) = X$. This auction is strategically equivalent to an $(n-1)$ -bidder standard FPSBA auction when the $n-1$ bidders face a seller that sets a naive reservation price strategy of $RP(X) = X$ as described in [2.9.4](#). As in the standard FPSBA derivation, this solution produces a non-linear term when the prior distribution has a variance $\sigma^2 \equiv \infty$.

Each bidder $i = 1, \dots, (n-1)$ observes a random signal X_i with variance σ_M^2 , which is an estimate of the unknown value V of the item being auctioned. $F(X_i|V)$ represents the conditional distribution of the i^{th} bidder's estimate, and $g(V|X_i)$ represents the bidder's posterior density for V given the signal X_i . Similarly, bidder n observes a random signal X_n with variance σ_T^2 . Let $B(X)$ represent the symmetric equilibrium strategy used by each of $n-1$ bidders, with $B'(\cdot) > 0$ such that the inverse $r(B) = B^{-1}$ is well defined. If all bidders $i = 1, \dots, n-1$ are using this strategy, and we now also assume that the n^{th} bidder adopts a naive bidding strategy that is fixed as $B_n(X) = X$ then the i^{th} bidder's problem may be written as:

$$\max_{B(X)} \int_{\Sigma V|X} (V - B) F_{\sigma_M}(r(B)|V)^{n-2} F_{\sigma_T}(r(X)|V) g(V|X) dV \quad (2.15)$$

Taking the first order condition $B'(X) = 0$ gives:

$$0 = \int_{\Sigma V|X} F_{\sigma_T}(r(X)|V) [(V - B)(n-2) F_{\sigma_M}(r(B)|V)^{(n-3)} f_{\sigma_M}(r(B)|V) r'(B) - F_{\sigma_M}(r(B|V)^{n-2}] g(V|X) dV \quad (2.16)$$

At the symmetric equilibrium, $r(B) = X$ and $r'(B) = B'(X)^{-1}$ and so :

$$B'(X) = \frac{\int_{\Sigma_{V|X}} (V - B) F_{\sigma_T}(r(X)|V) (n - 2) F_{\sigma_M}(X|V)^{n-3} f_{\sigma_M}(X|V) g(V|X) dV}{\int_{\Sigma_{V|X}} F_{\sigma_T}(r(X)|V) F_{\sigma_M}(X|V)^{n-2} g(V|X) dV} \quad (2.17)$$

A symmetric equilibrium must satisfy this differential equation 2.17. We now apply the three Thiel conditions as in section 2.9.3, making the substitution $z = X - V$, and additionally noting that the denominator is simply the probability of bidder i having the greatest signal denote by $\frac{1}{\alpha} (\neq \frac{1}{n})$:

Giving :

$$B'(X) = \frac{\int_{\Sigma_z} (X - B - z) (n - 2) F_{\sigma_T}(r(z)) F_{\sigma_T}(z)^{n-3} f_{\sigma_M}^2(z) dz}{\int_{\Sigma_z} F_{\sigma_T}(r(z)) F_{\sigma_M}(z)^{n-2} f_{\sigma_M}(z) dz (= \frac{1}{\alpha})} \quad (2.18)$$

$$\implies B'(X) + K_1 B(X) + K_2 - K_1 X = 0 \quad (2.19)$$

where

$$K_1 = \alpha \int_{\Sigma_z} (n - 2) F_{\sigma_T}(r(z)) F_{\sigma_M}(z)^{n-3} f_{\sigma_M}^2(z) dz \quad (2.20)$$

$$K_2 = \alpha \int_{\Sigma_z} z (n - 2) F_{\sigma_T}(r(z)) F_{\sigma_M}(z)^{n-3} f_{\sigma_M}^2(z) dz \quad (2.21)$$

Equation 2.19 is an ordinary linear differential equation in standard form, with solution:

$$B(X) = X - \frac{(1 + K_2)}{K_1} + \beta e^{-K_1 X} \quad (2.22)$$

Where β characterizes a family of solutions which are a consequence of the integration constant in 2.19, and we focus on the solution $\beta = 0$, (which is the limiting solution for the uninformative prior beliefs case with independently normal distribution of signals in the standard FPSBA). Equation 2.22 with $\beta = 0$ gives an equilibrium bidding function that is a linear function of the estimates, $B(X) = X - \delta$

Now applying the normally distributed signal condition gives the coefficients K_1 and K_2 :

$$K_1 = \alpha \int_{\Sigma_z} (n-2) F_{\sigma_T}(r(z)) \Phi_{\sigma_M}(z)^{n-3} \phi_{\sigma_M}^2(z) dz \quad (2.23)$$

$$K_2 = \alpha \int_{\Sigma_z} z(n-2) F_{\sigma_T}(r(z)) \Phi_{\sigma_M}(z)^{n-3} \phi_{\sigma_M}^2(z) dz \quad (2.24)$$

Where $\phi_{\sigma}(\cdot)$ and $\Phi_{\sigma}(\cdot)$ are the density and distribution function of a standard normal with variance σ^2 . Given equation 2.22, $B^{-1}(X) = X + \delta$, therefore, $z + \delta$ is also a normal RV and so $F_{\sigma_T}(r(z))$ is a normal CDF, which leads to:

$$F_{\sigma_T}(B^{-1}(z)) = \Phi_{\sigma_T}(z + \delta) = \Phi\left(\frac{z + \delta}{\sigma_T}\right) \quad (2.25)$$

Since z is $\sim N(0, \sigma_T)$ after the variable change in equation 2.18. Integrating 2.23 and 2.24 by parts and noting the order statistics result that the expected value of the maximum of n independent normal $Y = \max\{X_1, \dots, X_n\}$ but not identical distributions $X_i \sim N(\delta_i, \sigma_i)$ $i = 1, \dots, n$ can be written as :

$$E[Y] = \sum_{j=1}^n \frac{1}{\sigma_j} \int_{-\infty}^{\infty} s \phi\left(\frac{s - \delta_j}{\sigma_j}\right) \prod_{k=1, k \neq j}^n \Phi\left(\frac{s - \delta_k}{\sigma_k}\right) ds \quad (2.26)$$

Which is just the first moment of the first derivative of the distribution function of the maximum order statistic, $F = \prod_{k=1}^n \Phi\left(\frac{s - \delta_k}{\sigma_k}\right)$. Therefore, δ is the value that satisfies:

$$\delta = \frac{\sum_{j=1}^{n-1} \int_{-\infty}^{\infty} s^2 \phi\left(\frac{s}{\sigma_M}\right) \prod_{i=1}^{n-2} \Phi\left(\frac{s}{\sigma_M}\right) \Phi\left(\frac{s + \delta}{\sigma_T}\right) ds}{\sum_{j=1}^{n-1} \int_{-\infty}^{\infty} s \phi\left(\frac{s}{\sigma_M}\right) \prod_{i=1}^{n-2} \Phi\left(\frac{s}{\sigma_M}\right) \Phi\left(\frac{s + \delta}{\sigma_T}\right) ds} \quad (2.27)$$

Which, not surprisingly, is similar to the ratio of the 2nd to 1st moment solution in Wilson (1992) for the standard SBFPA normal case with $n-1$ bidders.

2.9.7 Summary

The equilibrium bidding solution in an n -bidder FPSBA⁴ where bidder $i = 1, \dots, n-1$ receive a signal $X_i \sim N(V, \sigma_M)$ and bidder- n receives a signal $X_n \sim N(V, \sigma_T)$ and

⁴Using the same methodology as above, can show that in an n -bidder auction where k -bidders adopt a naive strategy, the equilibrium bidding strategy for the $(n-k)$ strategic bidders is:

$$B(X) = X - \delta$$

adopts the naive strategy $B(X) = X$ is:

$$B(X) = X - \delta \quad (2.28)$$

Where δ satisfies:

$$\frac{\int_{-\infty}^{\infty} s^2 \phi\left(\frac{s}{\sigma_M}\right) \Phi^{n-2}\left(\frac{s}{\sigma_M}\right) \Phi\left(\frac{s+\delta}{\sigma_T}\right) ds}{\int_{-\infty}^{\infty} s \phi\left(\frac{s}{\sigma_M}\right) \Phi^{n-2}\left(\frac{s}{\sigma_M}\right) \Phi\left(\frac{s+\delta}{\sigma_T}\right) ds} - \delta = 0 \quad (2.29)$$

This is strategically equivalent to an $(n-1)$ bidder FPSBA where the seller adopts a naive reservation price $RP(X) = X$.

This function for δ (the amount of bid shaving) is not dependent on any signal realizations and is the value of δ that satisfies 2.29 and can only be solved numerically. This theoretical result, along with the empirical evidence of Cammack (1991) and Paarsch (1992) gives justification for examining a linear bidding strategy in the rest of the thesis.

2.9.8 Numerical Checks

As a numerical check, We examined three bidders (2 strategic 1 naive) and generated a vector of (1000×3) random normal samples, then calculated the rest response signal shaving given the other bidders' strategies. Figure 2.6 illustrates the result that the mutual best response point (the Nash equilibrium) coincides closely with the theoretical numerical solution of equation 2.29 as 1.73.

As a further numerical test, we attempted to reproduce the equilibrium spread levels with varying amounts of client signal variance that we previously calculated by numerically solving the payoff functions using the relaxation algorithm. The results appear to be identical and are shown in 2.7 for the 2, 3, 4, 5, and 10 bidder cases.

Where δ satisfies:

$$\frac{\int_{-\infty}^{\infty} s^2 \phi\left(\frac{s}{\sigma_M}\right) \Phi^{n-k}\left(\frac{s}{\sigma_M}\right) \Phi^k\left(\frac{s+\delta}{\sigma_T}\right) ds}{\int_{-\infty}^{\infty} s \phi\left(\frac{s}{\sigma_M}\right) \Phi^{n-k}\left(\frac{s}{\sigma_M}\right) \Phi^k\left(\frac{s+\delta}{\sigma_T}\right) ds} - \delta = 0$$





FIGURE 2.6: Best response functions of the 2 strategic bidders

FIGURE 2.7: Optimal bid reductions for the n strategic bidders with 1 naive bidder wrt relative signal variance

Chapter 3

Duopoly OTC Markets

3.1 Introduction

To analyze the equilibriums in a two-dealer OTC game, as defined in the previous chapter, we need the equilibrium results for a two-player Bertrand competition when there are unknown common values with asymmetric information. Although much has been written on Bertrand competition, there appears to be little in the competition literature concerning unknown common values with asymmetric information. The literature focus is primarily on an Independent Private Values modeling environment, where each firm or bidder has a private valuation or private marginal cost function, leading to a payoff that is the difference between the traded price and this private cost. The cases covered include where either market costs are known to all (the full information model, [Bertrand \(1883\)](#) and many subsequent complications) or where one-sided costs are unknown (partial information models, [Spulber \(1995\)](#), [Wambach \(1999\)](#), [Janssen and Rasmusen \(2001\)](#), [Patra et al. \(2019\)](#)).

The closest solution to this problem is [Spulber \(1995\)](#), who analyses an extension of the standard Bertrand model using an auction type representation in which marginal costs are not common knowledge among the competitors but are, for each firm, independently drawn from a distribution and private information for the individual firms. This is a one-sided information problem, where private costs are known but competitor's costs are unknown. This paper also uses an independent private value model and does not calculate the equilibrium's explicit form. This incomplete information about costs changes the equilibrium prediction dramatically, and prices are now set substantially above marginal costs. Other notable contributions are [Janssen](#)

and Rasmusen (2001), who looked at uncertainty in the number of competitors in price competition and found that prices rise with the reducing probability of greater numbers of competitors. Wambach (1999) introduced risk aversion into cost uncertainty, to name but a few, but these papers, in common with the majority, all model an Independent Private Values case, indeed according to Routledge and Edwards (2020), even the literature concerning informational uncertainty in competition is not large.

In the first part of this chapter, we examine the simple case of a common marginal cost to both firms (a common value model), combined with a limited information set regarding the true value of this cost, where only the estimated cost variance is known ante to setting a price. The common value case is a more realistic representation for markets where there is an unknown common cost element, for example, the amount of oil in a tract of land, the future raw materials cost in production. In particular, where there is an unknown common value to a financial instrument. However, the calculation of this value is subject to some measurement error or interpretation.

Price competition models play a significant role in public policy debates in regulation and antitrust; however, the standard full information Bertrand model of competition is subject to what is known as the Bertrand Paradox. This paradox is where two firms reach a Nash equilibrium state where both firms charge a price equal to marginal cost, share the market somehow, and earn zero profits.¹ However, the full information Bertrand model's standard conclusions critically depend on the assumption that firms have full information about the entire market cost structure.

We show that in the uncertain common cost Duopoly, each firm's uncertainty about the common marginal costs eliminates the discontinuity in the pricing strategy's effect on profit. Like Spulber (1995), we also show that all firms earn positive expected profits in equilibrium and are therefore incentivized to enter the market and engage in price competition. These results then translate directly into the bid-ask spread of dealers in an OTC market. The result of non-zero payoffs in equilibrium of a partial information Bertrand competition model is particularly relevant to other market models that assume a zero profit condition.

¹If all firms have constant marginal cost and if one firm has an absolute cost advantage over its rivals, it prices at the marginal cost of the next to lowest cost firm and captures the entire market. Thus, all other firms earn zero profit,

We model the competitive Bertrand price process as a sealed bid first price common value auction and use [Wilson \(1969\)](#) standard equilibrium results for a two player auction. (whose results were extended in the seminal paper by [Milgrom and Weber \(1982\)](#), who provided a full if a somewhat unfriendly description of the general form of the equilibrium in auctions.)

This chapter also looks specifically at the case of asymmetric information between the two firms in a duopoly. We show that in an uncertain common cost environment, as uncertainty regarding the true value goes up for either firm, both firms' equilibrium price also goes up. This is due to the firms' strategic nature when they know the level of uncertainty (variances) in the market, and the level of uncertainty is common knowledge. Intuitively, if one knows one's competitor is very unsure about the true costs, one can be less aggressive in one's own pricing strategy. This result gives further evidence that public information can reduce spreads in the market. It has the biggest effect on the worst-informed firms, which affects the best-informed firms' ability to exploit their informational advantage.

We also look at asymmetric information between the dealers and the clients and demonstrate how dealers take advantage of their informational advantage. This result is particularly relevant to OTC markets since clients are often at an informational disadvantage to the dealers.

According to [Routledge and Edwards \(2020\)](#), "our understanding of price competition in the presence of production cost uncertainty is still rudimentary", notwithstanding the large volume of work on competition spawned by Bertrand and Cournot. Even concerning the century-old standard Bertrand model, "there is a notable gap in the research. There are no equilibrium existence results for the classical Bertrand model when there is discrete cost uncertainty."

This section adds to the literature in two directions. Firstly, examining in detail the equilibrium in a Bertrand Duopoly price competition with unknown common costs and asymmetric information using an auction methodology. Secondly, in applying this competition methodology to financial market instruments, where market makers replace the traditional firms and marginal cost is replaced by a 'true' common value, and pricing takes the form of a bid-ask spread. Bertrand competition's pricing

equilibria then translate into predictions for the equilibrium bid-ask spread in financial markets, which typically accords with the main microstructure models' findings but allows for greater flexibility when applied to a network and dispenses with some empirically questionable assumptions.

The extent of the firms' information sets when making a price is critical to which price they make. In standard Bertrand price competition literature, the true cost is the marginal cost to the firm of supplying the object. In contrast, in financial markets, the true cost is the true value but is realized only after trading. Agents attempt to measure this true cost, and we model this by them receiving a sample (signal) from a normal distribution with an unknown mean but known variance. The assessment of the true value by the firms of the marginal cost or value is subject to measurement error or interpretation.

3.2 The Model

We begin by looking at the dual uncertain information duopoly price problem, a standard Bertrand competition model with unknown common costs to both players: two identical firms who compete in price for a homogeneous good. Each firm has a common marginal cost C , which is unknown to both firms a priori. We simplify the market structure by assuming a unit demand from each customer; however, the results from a linear demand curve are shown to have very similar effects. Next, we apply this common value auction methodology to financial market prices where firms provide a firm buy and sell price to the market customers. The firm with the lowest price matches the customer buy order, and the highest buy price matches the customer sell order. The payoff to the firms is price - cost, where the price is the best-quoted trading price and the cost is the future true value (or can be interpreted as the future mark to market value). We can model this price competition (Bertrand) as a common value first price sealed bid auction where the object's true cost is unknown but the same to each bidder. In the auction setup, the customer acts as the auctioneer and the firms as the bidders.

3.2.1 Information Sets

The uncertainty of the firms regarding the true cost C defines a game of incomplete information. In line with the methodology of [Harsanyi \(1967\)](#), we convert this game to a game of imperfect information by assuming a probability distribution over the true value. As each firm does not know the other firm's signal, each firm forms a probability distribution belief around competitor types (error signal distributions), which we assume is common knowledge along with the bidding functions and own signal error distributions. This is the standard technique, known as the Common Prior Assumption, for converting games of incomplete information to games of imperfect information. We assume that the distribution of each firm's error signal is common knowledge which is itself common knowledge. This allows us to specify this structure as a static Bayesian game and the Bayesian Nash Equilibrium as the equilibrium concept.

Specifically, we assume that the true value estimation is modeled as a draw from a probability distribution. The firms know the functional form of the distribution (normal) and know the variance. However, the mean (the true value) is unknown.

3.2.2 Normal Distribution of the Estimation (Signal) Errors

The extent of the information sets that firms have when making a price is critical in what price they make. Our motivation in this chapter is that firms in general and financial market participants, in particular, have a payoff that is essentially (price traded - cost), where the true cost is unknown. In financial markets, the true cost can be replaced with the true value since the financial instrument will realize a true value (bond repayment, derivative expiry etc.) at some future time. As in [Chapter 2](#), agents attempt to measure this true cost (value), and we can model this by each of them receiving a sample (signal) from a normal distribution.

Therefore, we model the true value estimation process by each firm i of common cost C by them receiving an independent signal S_i where firm i receives a signal $S_i = C + \epsilon_i$ with ϵ_i normal $N(0, \sigma_i)$.

3.3 Optimal Bidding Strategies

3.3.1 The Standard Auction Model

The most common equilibrium bidding model is a non-cooperative Nash equilibrium with risk-neutral bidders, known as a risk-neutral Nash equilibrium (RNNE) solution. [Wilson \(1969\)](#) was the first to develop an RNNE solution to the common value auction problem, and this methodology was significantly extended to include private values by the seminal papers of [Milgrom and Weber \(1982\)](#).

Although auctions of varying types have been analyzed for a very long time, Milgrom and Weber's paper provides a comprehensive solution to the equilibrium of various auction formats, which launched many papers with varying degrees of complication. However, in general, the solutions to the equilibriums in Milgrom equations are not simple to explicitly find. [Athey \(1997\)](#), [Lizzeri & Persico \(1998\)](#) and [Rodriguez \(2000\)](#) prove existence and uniqueness results. Athey also develops numerical algorithms for computing the equilibrium. [Laskowski & Slonim \(1999\)](#) provide an asymptotic solution for a parametric model. [Kagel & Levin \(1999\)](#) exhibit a bounded rational solution for a similar model. [Hausch \(1987\)](#) determines the equilibrium for a discrete setting. [Campbell & Levin \(2000\)](#) solve the equilibrium for a discrete, parametric model.

The Wilson Model

Wilson set up the framework for the equilibrium solution of the common value problem. The following brief summary is taken from his seminal paper. Suppose that two parties, called 1 and 2, will bid for a prize of monetary value v which is not known with certainty by either party. For simplicity, assume that both parties initially assess the same prior probability density, $g(v)$, for the value of the prize v . Then, before the bidding, each party i observes an outcome θ_i of a random variable $\hat{\theta}_i$ distributed with the conditional density $h_i(\theta_i|v)$. We assume that conditional on v , $\hat{\theta}_1$, and $\hat{\theta}_2$ are independent.

Our aim is to identify the equilibrium pure strategies when they exist; say, $p_i(\theta_i)$ is the bid to be made by party i if he observes θ_i . He then shows that a pure

strategy function $p_i(\theta_i)$ is monotonic and so an inverse function $\Pi_i(p_i)$ exists satisfying $\Pi_i(p_i(\theta_i)) = \theta_i$. He further assumed that each inverse function Π_i is differentiable.

Finally, a utility function for money that is linear in money is assumed for each party. Each agent chooses his bidding strategy to maximize his expected net gain if he should win, given his opponent's strategy. Suppose party 2 chooses $p_2(\theta_2)$ as a strategy.

The final solution to the equilibrium is the solution of a differential equation of the form:

$$\frac{d\Pi_2}{d\Pi_1} = \frac{\varphi_{21}(\Pi_2|\Pi_1)}{\varphi_{21}(\Pi_1|\Pi_2)} \quad (3.1)$$

where, $f_{ji}(\theta_j|\theta_i) = \int_{-\infty}^{\infty} h_j(\theta_j|v)g_i(v|\theta_i) dv$ is the posterior marginal density for his opponent's observation, and the posterior marginal distribution function is $F_{ji}(\theta_j|\theta_i) = \int_{-\infty}^{\theta_j} f_{ji}(\xi|\theta_i) d\xi$ and the function $\varphi_{ji}(\theta_j|\theta_i) := \frac{F_{ji}(\theta_j|\theta_i)}{f_{ji}(\theta_j|\theta_i)}$ which yields Π_1 as a function of Π_2 .

Let $\bar{v}(\theta_1, \theta_2)$ be the expected value of the common value v , conditional on the observations θ_1 and θ_2 , then;

$$\bar{v}(\theta_1, \theta_2) = \int_{-\infty}^{\infty} v[h_2(\theta_2|v)g_1(v|\theta_1)/f_{21}(\theta_2|\theta_1)] dv \quad (3.2)$$

and the equilibrium solution is obtained by solving the following equation that is derived from the first order condition:

$$\varphi_{ji}(\Pi_j(p)|\Pi_i(p)) = [\bar{v}(\Pi_1(p), \Pi_2(p)) - p]\Pi_j'(p) \quad (3.3)$$

In the independent information case, $\bar{v}(\Pi_1(p), \Pi_2(p)) = \Pi_1$ and by inspection of the partial differential equation PDE (3.3), a linear solution is apparent. Consider a candidate linear solution, i.e, $p = \Pi_i + \beta$, for some β , then $\Pi_i'(p) = 1$ and its easy to check that this a solution to the PDE when $\beta = \varphi(\Pi_j(p)|\Pi_i(p))$, and so $p = \Pi_i - \beta$.

Suppose that the two parties have the same types of information available, meaning that h_1 and h_2 are identical functions and, therefore, that φ_{12} and φ_{21} are identical functions. In this case, a solution to (3.1) is $\Pi_2(p) = \Pi_1(p)$. Further, if the common prior assessment g is a diffuse Normal density and the observations θ_i are each

Normally distributed with mean v , then each posterior marginal density f_{ji} is a Normal density function with mean Π_i and a common standard deviation, say σ ; hence, $F_{ji}(\Pi_1|\Pi_1) = \frac{1}{2}$, $f_{ji}(\Pi_1|\Pi_1) = \frac{1}{\sigma\sqrt{2\pi}}$ and $\varphi_{ji}(\Pi_1|\Pi_1) = \sigma\sqrt{\frac{\pi}{2}}$ $\bar{v}(\Pi_1, \Pi_1) = \Pi_1$. The solution to PDE (3.3) is therefore $\Pi_1(p) = \Pi_2(p) = p + \sigma\sqrt{\frac{\pi}{2}}$ and the optimal strategy functions are $p_i(\theta_i) = \theta_i - \sigma\sqrt{\frac{\pi}{2}}$, $i = 1, 2$.

Therefore, the equilibrium strategy is calculated as $P_i(S_i) = S_i - \sigma\sqrt{\frac{\pi}{2}}$, where σ^2 is the variance of the posterior normal distribution (after observing your sample). From Bayesian statistical updating, the prior standard deviation, σ_{prior} , and the posterior deviation, σ_{post} , after 1 sample are related by:

$$\sigma_{post} = \left(\frac{1}{\sigma_{prior}^2} + \frac{1}{\sigma_{prior}^2} \right)^{-1} = \frac{\sigma_{prior}}{\sqrt{2}} \quad (3.4)$$

Therefore, in relation to the prior variance σ of the error distribution (which we assume is common knowledge), the Nash equilibrium as calculated by Wilson can be rewritten as:

$$P_i(S_i) = S_i - \sigma\sqrt{\pi} \quad (3.5)$$

3.3.2 Linear Bidding Functions

As we are concerned with a pure common value representation, we use the results of [Wilson \(1969\)](#), who gave a linear solution to the simple normal distribution error case $N(0, \sigma)$ in a common value sealed bid auction with unit demand function and two bidders. The solution is of the form $b(S_i) = S_i + \sigma\sqrt{\pi}$, and was derived by solving the bidding function as a solution to first-order conditions of the payoff functions leading to a differential equation and certain boundary conditions.

Section 2.9 gives a theoretical basis for examining a linear bidding function in the more general n -bidder case with normal signal error. Empirically, [Cammack \(1991\)](#) ran multiple regressions on the Treasury Bill auction dataset, where traders placed bids in a Treasury market auction, and she finds: "They imply a linear functional form between the bidding adjustment and the empirical measures of dispersion of opinion and number of bidders."

This theoretical and empirical evidence motivates us to examine a symmetric linear bidding function for the firms. Specifically, when a firm receives an estimate

S_i , they are aware that if they bid their estimate, this is likely to be too high and leads to a negative expected payoff. This is because their estimate of the common value is the highest amongst the other competing firms, and on average, will give a negative payoff. This is because the probability that the estimate is too high, conditional on it being a maximum, is greater than the unconditional probability that the estimate is too high. This is known as the 'Winners Curse'. See [Kagel and Levin \(1986\)](#) for a comprehensive study into this effect. The firms are strategic in the sense that they are aware of this conditional probability problem and therefore seek to shave an amount from their estimate.

3.4 Standard Duopoly Market

3.4.1 Assumptions

Consider a standard duopoly market with a single customer and two firms that compete in price for an identical good. The customer has unit demand and trades at the best price of the two firms that he is connected to. Each firm can produce this homogenous good that depends entirely on an uncertain common cost C .

Each firm i attempts to estimate this common cost C and receives an independent signal S_i where firm 1 receives a signal $S_1 = C + \epsilon_1$ with ϵ_1 normal $N(0, \sigma_1)$ and firm 2 receives a signal $S_2 = C + \epsilon_2$, where ϵ_2 is $N(0, \sigma_2)$ to allow for differences in estimation precision.

Each firm adds a profit margin δ_i to their signal and each seek to maximize their payoff. Firm 1 makes a price $P_1(S_1) = S_1 + \delta_1$ and firm 2 makes a price $P_2(S_2) = S_2 + \delta_2$. The firm with the lowest price gets to trade and has a payoff of $P_i - C$ or zero if the price is higher than the other firm.

We make the common assumption that the firms are risk-neutral and seek to maximize their expected profits in equilibrium.

This setup can lead to a winner's curse problem - if δ_i is too low, then the probability of trading is increased, but expected profits are low (possibly negative), and if δ_i is too high, the probability of winning is low but expected profitability is high. The Winner's Curse is the name of the phenomena that is the tendency for the winning bid in an auction to exceed the intrinsic value or true worth of an item.

3.4.2 Game setup

From a game-theoretic perspective, we can model this partial information competitive process as a Bayesian game. The Bertrand game follows the usual protocols. Firms 1, 2, receive signals S_i as to the true marginal common cost C . Firms make a price $P_i(S_i)$, where the function $P_i(\cdot)$ is known as a bidding function and is a strategy that maps each possible cost estimation onto a trading price. The firms do not see the other firms' price before making their own price. The client then selects the firm with the lowest price to trade. The true cost C is realized, and the payoffs are $P - C$ for the best-priced firm, which captures the whole market. The payoff is zero for the other firm.

Players: Firms 1 and 2

Typesets: Each player receives a signal S_i drawn from a distribution $N(C, \sigma)$ - the estimate of the cost C determines types

Strategies: Each firm i , makes a price P , with a bidding function $P(S_i) = S_i + \delta_i$, $\delta_i \in [0, \infty)$, $i = 1, 2$, this strategy maps each possible cost estimate onto a trading price

Payoffs: Payoffs are $P_i - C$ if $P_i < P_j$, zero otherwise

3.4.3 Expected Payoffs

Each firm $i=1,2$ makes a selling price (after receiving their signal S_i) of

$$P(S_i) = S_i + \delta_i = C + \varepsilon_i + \delta_i \quad (3.6)$$

and has a payoff of

$$\Pi_i = (P(S_i) - C)\mathbb{1}_{(\varepsilon_i + \delta_i) < (\varepsilon_j + \delta_j)} = (\varepsilon_i + \delta_i)\mathbb{1}_A \quad (3.7)$$

where $\mathbb{1}_A$ is an indicator function that is equal to 1 if $(\varepsilon_i + \delta_i) < (\varepsilon_j + \delta_j)$ of the other firm, zero otherwise.

We are interested in the probability that firm i 's price is less than firm j 's price. Formally we are looking for the probability of firm i , to win a competitive auction with one other firm, that is;

$$P(P_i < P_j) = P(V + \varepsilon_i + \delta_i < V + \varepsilon_j + \delta_j) \quad (3.8)$$

Consider $N+1$ independent normal random variables, X_0, X_1, \dots, X_N , where each X_i has a mean δ_i and standard deviation σ_i . It is standard theory² that the probability of an independent normal RV being less than the minimum of N other normal RV's is

$$\Pr(X_0 < Y) = \int_{-\infty}^{\infty} \frac{1}{\sigma_0} \phi\left(\frac{s - \delta_0}{\sigma_0}\right) \prod_{j=1}^N \left(1 - \Phi\left(\frac{s - \delta_j}{\sigma_j}\right)\right) ds \quad (3.9)$$

And the expected value of this RV is:

$$E[X_0 : X_0 < Y] = \int_{-\infty}^{\infty} s \Pr(X_0 = s) \Pr(Y > s) ds \quad (3.10)$$

Therefore, using this result allows us to specify the expected payoff to firm i as:

$$E[\Pi_i] = E[(\varepsilon_i + \delta_i) \mathbf{1}_A] = \int_{-\infty}^{\infty} \frac{s}{\sigma_{Mi}} \phi\left(\frac{s - \delta_i}{\sigma_{Mi}}\right) \left(1 - \Phi\left(\frac{s - \delta_j}{\sigma_{Mj}}\right)\right) ds \quad (3.11)$$

This formulation can be easily extended to n firms, but this integral (3.11) has an analytic solution only in the 2 firm case, and the Expected payoff of firm 1, given $\delta_1, \delta_2, \sigma_1, \sigma_2$: can be shown to be:

$$E[\Pi_1(\delta_1, \delta_2, \sigma_1, \sigma_2)] = \delta_1 \Phi\left(\frac{\delta_2 - \delta_1}{\theta}\right) - \frac{\sigma_1^2}{\theta} \phi\left(\frac{\delta_2 - \delta_1}{\theta}\right) \quad (3.12)$$

where $\theta = \sqrt{\sigma_1^2 + \sigma_2^2}$ and $\Phi(\cdot)$ and $\phi(\cdot)$ are the distribution and density functions of a standard normal distribution $N(0, 1)$.

Similarly firm 2's expected payoff is:

$$E[\Pi_2(\delta_1, \delta_2, \sigma_1, \sigma_2)] = \delta_2 \Phi\left(\frac{\delta_1 - \delta_2}{\theta}\right) - \frac{\sigma_2^2}{\theta} \phi\left(\frac{\delta_1 - \delta_2}{\theta}\right) \quad (3.13)$$

These expected payoff functions Eqn (3.12) and Eqn (3.13) are decreasing as a function of σ_1 for firm 1 and similarly decreasing in σ_2 for firm 2. Superior information (estimation of true value) translates into higher expected payoffs.

²See for example, Hill (2010) - Minimum of Normally distributed random variables

As both variances tend to zero $\lim_{\sigma_i \rightarrow 0} E[\Pi_1] = \delta_1$ if $\delta_1 < \delta_2$, $\frac{\delta_1}{2}$ if $\delta_1 = \delta_2$, 0 if $\delta_1 > \delta_2$, which is the full information set normal Bertrand competition payoffs with known costs with the familiar discontinuity and the cause of the leapfrogging by each firm to a price of marginal cost.

3.4.4 Equilibrium

Symmetric Information Case

This 2-player non-cooperative game has a Bayesian Nash Equilibrium by standard existence theorems and a pure strategy risk neutral (Bayesian) Nash equilibrium was calculated analytically in the case of a traditional auction for the common estimation error case, $\sigma_1 = \sigma_2 = \sigma$ by Wilson (1969, 1977) using the solution to the first order conditions of a general bidding function. Later, again using an auction setup, Thompson et al. (2005) assumed a linear type bidding function and solved the first order conditions of the maximum order statistics using several standard probability distributions. They both lead to the same result as:

$$\delta^*(\sigma) = \frac{\frac{1}{2} + \int_{-\infty}^{\infty} x \phi_{\sigma}^2(x) dx}{\int_{-\infty}^{\infty} \phi_{\sigma}^2(x) dx} = \sigma \sqrt{\pi} \quad (3.14)$$

Where ϕ_{σ} is the normal density function with variance σ^2 which is the variance of the distribution of the normal errors.

The expected payoff at this equilibrium is :

$$E[\Pi_i(\delta^*, \delta^*, \sigma, \sigma)] = \sigma \left(\frac{\pi - 1}{2\sqrt{\pi}} \right) \quad (3.15)$$

Interestingly, the payoff for both firms at this equilibrium is positive for non-zero σ - in contrast to the full information Bertrand models. This is similar to the results of Spulber, who found that uncertainty over competitor's costs created non-zero profits in equilibrium (although this was in an Independent Private Values setting). This non-zero profit in equilibrium in the presence of common cost uncertainty is another solution to the classic Bertrand Paradox. When $\sigma \rightarrow 0$, $\delta^* \rightarrow 0$ and $E[\Pi] \rightarrow 0$ as in standard full information Bertrand model.

The more general problem with a greater numbers of firms can be solved in a similar way but unfortunately the solutions, which can only be expressed as complicated integrals, can only be solved numerically and offer less insight into the nature of the equilibrium. However, we provide an analysis of the equilibrium in the important case of asymmetric information in the 2 firm case below.




FIGURE 3.1: Theoretical and Numerical Simulation of Signal Reduction Best Response Curves - Firm1(blue) $N(0,1)$, Firm2(red) $N(0,1)$ signal error.

To gain intuition, Figure 3.1 illustrates the best response functions of the 2 firms who both receive a $N(0, 1)$ error signal, and compares the theoretical predictions with numerical simulations of the mean payoffs from 10,000 auctions. At the intersection point, both firms have a mutual best response to each others' strategy and hence is a Nash equilibrium. The numerical intersection is seen to be consistent with the theoretical result - $\sqrt{\pi}$ approx 1.7725.

Asymmetric Information

An explicit expression for the Nash equilibrium exists in the uncertain and asymmetric information case, but simple best response functions cannot be expressed explicitly and must be calculated numerically. For each choice of firm j , we calculate the best response of firm i . We do the same for firm i and find the point of intersection of these best response curves. For intuition, consider the Duopoly market as mentioned above but now with firm 2 having a less accurate ability to measure the true common cost. This situation can be modeled by having firm 2 receive a signal that has a higher variance than firm 1.

Figure 3.2 graphically illustrates the best response functions of the two firms when firm 1 has a signal error of $N(0, 1)$ and firm 2 has a signal error of $N(0, 2)$. The intersection of the best response curves is the Nash equilibrium of the system as




FIGURE 3.2: Signal Reduction Best Response Curves - Firm 1 has $N(0,1)$ and Firm 2 has $N(0,2)$ signal error

this is a mutual best response to each other's strategies. The best response function of firm 1 (the low variance firm) is less sensitive to firm 2's strategy than vice versa. In fact, numerical simulations suggest that the equilibrium levels are equal for any combination of σ_1 and σ_2 . The payoffs for the 2 firms at these levels are, of course, very different, with the lower signal error firm making significantly higher payoffs. This motivates the following proposition which allows us to explicitly calculate the equilibrium level.

Proposition 1. *Given the auction duopoly with unknown common values described above and firms receive signals $S_1 \sim N(C, \sigma_1)$ and $S_2 \sim N(C, \sigma_2)$ regarding the true cost, then in equilibrium, where the 2 firms use a bidding strategy $B(S_i) = S_i + \delta_i^*$, then for any σ_1, σ_2 , we have $\delta_1^* = \delta_2^*$.*

Proof. Assume equilibrium and both firms are choosing their mutual best responses (δ_1^*, δ_2^*) . Consider firm 1, they have an expected payoff of

$$E[\Pi_1] = \delta_1 \Phi\left[\frac{-\delta_1 + \delta_2}{\sqrt{\sigma_1^2 + \sigma_2^2}}\right] - \frac{1}{\sqrt{2\pi}} \frac{\sigma_1^2}{\sqrt{\sigma_1^2 + \sigma_2^2}} e^{-\frac{(\delta_1 - \delta_2)^2}{2(\sigma_1^2 + \sigma_2^2)}}$$

Let $a = \sqrt{\frac{\pi}{2}}$, $Z = \sqrt{\sigma_1^2 + \sigma_2^2}$, $\delta_2^* = aZ$ and let $\delta_1^* = aZ + \Delta$, where Δ now becomes firm 1's choice variable.

$$E[\Pi_1] = (aZ + \Delta) \Phi\left[\frac{-\Delta}{\sqrt{Z}}\right] - \frac{1}{\sqrt{2\pi} \frac{\sigma_1^2}{Z}} e^{-\frac{\Delta^2}{Z^2}} \quad (3.16)$$

At the equilibrium, the choice of Δ is a best response of firm 2's choice, hence:

$$\frac{\delta E[\Pi_1]}{\delta \Delta} = \Phi\left(\frac{-\Delta}{Z}\right) + (1-a)\phi\left(\frac{\Delta}{Z}\right) + \frac{2\Delta}{C^2\sqrt{2\pi}}e^{\frac{-\Delta^2}{Z^2}} = 0$$

$\Delta = 0$ is clearly a solution and can be shown to be unique. Similarly, examining the payoffs of firm 2 and setting $\delta_2^* = \delta_1^* + \Delta'$ gives identical expressions which demonstrate that $\Delta' = \Delta = 0$, hence, $\delta_1^* = \delta_2^*$

□

This proposition has some interesting consequences. In equilibrium, the 2 firms tend to make the same signal reduction, independent of their estimation error. In fact, given that the equilibrium spreads are the same, the expected payoffs for the 2 firms are:

$$E[\Pi_1] = \frac{\delta^*}{2} - \frac{\sigma_1^2}{\sqrt{(\sigma_1^2 + \sigma_2^2)}}, E[\Pi_2] = \frac{\delta^*}{2} - \frac{\sigma_2^2}{\sqrt{(\sigma_1^2 + \sigma_2^2)}} \quad (3.17)$$

The payoffs are decreasing with standard deviation of the error signal and increasing in the uncertainty of the competitor.

In fact, given this proposition and the payoff functions, we can state the following:

Proposition 2. *Given the auction duopoly with unknown common values described above and firms receive signals $S_1 \sim N(C, \sigma_1)$ and $S_2 \sim N(C, \sigma_2)$ regarding the true cost, then in equilibrium, where the 2 firms use a bidding strategy $B(S_i) = S_i + \delta_i^*$, then for any σ_1, σ_2 , we have that both firms have a Nash equilibrium strategy $\delta_1^* = \delta_2^* = \delta^*$, where $\delta^* = \sqrt{\frac{\pi}{2}(\sigma_1^2 + \sigma_2^2)}$*

Proof. Follows from proposition 1. Set $\delta_2^* = \delta_1^* + z$, take partial derivative of payoff 2 wrt z and solve the first order condition which leads to an expression for δ_1^*

$$E[\Pi_2(\delta_1^*, \delta_2^*, \sigma_1, \sigma_2)] = (\delta_1 + z)\Phi\left[\frac{-z}{\sqrt{\sigma_1^2 + \sigma_2^2}}\right] - \frac{1}{\sqrt{2\pi}} \frac{\sigma_1^2}{\sqrt{\sigma_1^2 + \sigma_2^2}} e^{-\frac{(z)^2}{2(\sigma_1^2 + \sigma_2^2)}}$$

The first order condition wrt to z is:

$$E_z[\Pi_2] = \Phi\left(\frac{-z}{\sqrt{\sigma_1^2 + \sigma_2^2}}\right) - \frac{\delta_1^* + z}{\sqrt{\sigma_1^2 + \sigma_2^2}} \phi\left(\frac{-z}{\sqrt{\sigma_1^2 + \sigma_2^2}}\right) + \frac{2z}{\sigma_1^2 + \sigma_2^2} \frac{\sigma_1^2}{\sqrt{2\pi(\sigma_1^2 + \sigma_2^2)}} e^{-\frac{z^2}{2(\sigma_1^2 + \sigma_2^2)}}$$

At equilibrium $z=0$ by Proposition 1, therefore, after some rearranging:

$$\delta_1^* = \delta_2^* = \frac{\sqrt{(\sigma_1^2 + \sigma_2^2)}\Phi(0)}{\phi(0)} = \sqrt{\frac{\pi}{2}(\sigma_1^2 + \sigma_2^2)} \quad (3.18)$$

□

Also it is easy to check that this is an equilibrium, which then also proves proposition 1.

This expression demonstrates that as uncertainty regarding the true value goes up for either firm, the equilibrium price goes up. The increase in the equilibrium price is due to the firms' strategic nature when they know the level of uncertainty (variances) in the market. As the uncertainty goes up for one firm, they need to make greater signal reductions to protect themselves from the winner's curse. The other firm knows this and can afford a greater signal reduction, leading to a higher payoff. As usual, the full information case is covered with $\sigma_1 = \sigma_2 = 0$ and gives the standard result of marginal cost pricing in equilibrium.

The expected payoff at this equilibrium is :

$$E[\Pi_1] = \frac{\sqrt{\frac{\pi}{2}(\sigma_1^2 + \sigma_2^2)}}{2} - \frac{\sigma_1^2}{\sqrt{2\pi(\sigma_1^2 + \sigma_2^2)}} = \frac{\sigma_1^2(\pi - 2) + \pi\sigma_2^2}{\sqrt{8\pi(\sigma_1^2 + \sigma_2^2)}} \quad (3.19)$$

$$E[\Pi_2] = \frac{\sqrt{\frac{\pi}{2}(\sigma_1^2 + \sigma_2^2)}}{2} - \frac{\sigma_2^2}{\sqrt{2\pi(\sigma_1^2 + \sigma_2^2)}} = \frac{\sigma_2^2(\pi - 2) + \pi\sigma_1^2}{\sqrt{8\pi(\sigma_1^2 + \sigma_2^2)}} \quad (3.20)$$

These are simply derived by inserting the equilibrium values into the individual payoff functions (3.12), (3.13). These expected payoffs are always non-zero for any non-zero sum of the variances ($\sigma_1^2 + \sigma_2^2$), illustrating that even a less informed firm can compete profitably with a better informed firm in equilibrium. These results are critically dependent on the information sets of the 2 firms and this particular observation is a consequence of both firms knowing the variances, σ_1^2 and σ_2^2 of the other firm's signal error and pricing strategically.

The fact that in equilibrium, the prices rise for both firms if one of the firms is less well informed is analogous in some ways to the standard full information Bertrand model where one firm has a private cost advantage over the other firm. The Nash

equilibrium of this standard duopoly model is that the equilibrium price is the higher-cost firm's marginal cost.

There is some empirical evidence to suggest this effect in real pricing - the dispersion of bid-ask spreads in financial markets is relatively small considering the various firms' varying estimation abilities and myriad sources of measurement error. The consistent high profitability of the major financial firms suggests that their spreads are very profitable. In contrast, small market-making firms appear less so. [Huang and Masulis \(1999\)](#) analyzed the dealer bid-ask spreads in the foreign exchange market, which has some similarities to our stylized model. The fx market is an OTC market with no central marketplace, and clients contact dealers directly to ask for trading quotes. They found that average dealer spreads to customers were 0.807% with a standard deviation of 0.0137%

3.4.5 Customer Payoffs

The expected payoff $E[\Pi_C]$, to the customer is the negative sum of the payoffs of both firms, which is :

$$E[\Pi_C(\delta_1, \delta_2, \sigma_1, \sigma_2)] = \theta \phi \left(\frac{\delta_1 - \delta_2}{\theta} \right) - \delta_1 \Phi \left(\frac{\delta_2 - \delta_1}{\theta} \right) - \delta_2 \Phi \left(\frac{\delta_1 - \delta_2}{\theta} \right) \quad (3.21)$$

which is the standard formula for the expected value of the minimum of two normal random variables $N(\delta_1, \sigma_1)$, $N(\delta_2, \sigma_2)$, as in [Nadarajah and Kotz \(2008\)](#), who also extended this formula to include correlated error signals and if $\theta = \sqrt{\sigma_1^2 + \sigma_2^2 - 2\rho\sigma_1\sigma_2}$, where ρ is the correlation coefficient between the two samples (or estimation errors), the Eqn (3.21) is the more general expression for the expected payoff.

At equilibrium, $\delta_1^* = \delta_2^*$ and the expected payoff to the customer is :

$$E[\Pi_C(\delta_1, \delta_2, \sigma_1, \sigma_2)] = \sqrt{\frac{\sigma_1^2 + \sigma_2^2 - 2\rho\sigma_1\sigma_2}{2\pi}} - \delta^* \quad (3.22)$$

Equation (3.22) is decreasing in the correlation coefficient ρ , that is, as firms error signals become more correlated, expected payoffs to the customer decrease,

hence firms' payoffs increase. As error signals become increasingly correlated, the competition effect diminishes. This in turn causes equilibrium spreads to decrease.

3.4.6 Non Unit Demand Functions

So far the payoff functions have assumed a unit demand. The equilibriums calculated also hold for a classic downward sloping linear demand function in the symmetric information case. Suppose the duopoly is set up as above, with firm 1 receiving a signal $S_1 \sim N(0, \sigma_1)$ and firm 2 receives signal $S_2 \sim N(0, \sigma_2)$. Suppose that both firms face a market demand function $D(P) = aP + c$, $a < 0$. The payoff to firm 1 is:

$$\Pi_1 = D(P)(P - C) = D(P(S_1))(P(S_1) - C) = D(V + \epsilon_1 + \delta_1)(\epsilon_1 + \delta_1)$$

if $(\epsilon_1 + \delta_1)$ is less than $(\epsilon_2 + \delta_2)$, zero otherwise. Analogous expression for firm 2.

Then the expectation of this payoff $E[\Pi_1]$ is:

$$E[\Pi_1] = E[D(V + \epsilon_1 + \delta_1)]E[\epsilon_1 + \delta_1] - 2COV[D(V + \epsilon_1 + \delta_1), \epsilon_1 + \delta_1]$$

In the linear case $D(P) = aP + c$, this becomes:

$$E[\Pi_1] = E[a(V + \epsilon_1 + \delta_1) + c]E[\epsilon_1 + \delta_1] - 2Var[\epsilon_1 + \delta_1] \quad (3.23)$$

The variance of the minimum order statistic $(\epsilon_1 + \delta_1)$ can be calculated ³

$$Var[\epsilon_i + \delta_i] = \int_{-\infty}^{\infty} \frac{s^2}{\sigma_1} \phi\left(\frac{s - \delta_1}{\sigma_1}\right) \Phi\left(\frac{s - \delta_2}{\sigma_2}\right) ds - \left(\int_{-\infty}^{\infty} \frac{s}{\sigma_1} \phi\left(\frac{s - \delta_1}{\sigma_1}\right) \Phi\left(\frac{s - \delta_2}{\sigma_2}\right) ds \right)^2 \quad (3.24)$$

At equilibrium, $\delta_1 = \delta_1^*$ and hence maximizes $E[\epsilon_1 + \delta_1]$ (the unit demand case), and $D(x)$ is a decreasing function, therefore δ_1^* also maximizes $E[D(V + \epsilon_1 + \delta_1)]E[\epsilon_1 + \delta_1]$. In the symmetric information case δ_1^* must equal δ_2^* by symmetry, therefore, $Var(\epsilon_1 + \delta_1^*)$ is the same as $Var(\epsilon_1 + \delta_2^*)$. Therefore, at equilibrium, the Nash equilibrium pair (δ_1^*, δ_2^*) of the unit demand case, is also the same equilibrium of the linear demand

³The first two moments of $Y = \min(X_1, X_2)$ are $E[Y] = \mu_1 \Phi\left(\frac{\mu_2 - \mu_1}{\theta}\right) + \mu_2 \Phi\left(\frac{\mu_1 - \mu_2}{\theta}\right) - \theta \phi\left(\frac{\mu_2 - \mu_1}{\theta}\right)$ and $E[Y^2] = (\sigma_1^2 + \mu_1^2) \Phi\left(\frac{\mu_2 - \mu_1}{\theta}\right) + (\sigma_2^2 + \mu_2^2) \Phi\left(\frac{\mu_1 - \mu_2}{\theta}\right) - (\mu_1 + \mu_2) \theta \phi\left(\frac{\mu_2 - \mu_1}{\theta}\right)$, see Nadarajah (2008)

case. This allows us to restrict the analysis to the unit demand case as the algebra is simpler, however, all of the previous results hold for the more realistic case of a linear demand function.

In the asymmetric σ_i case, the equal equilibrium price levels no longer hold as the conditional variance (3.24) of the minimum order statistic is not symmetrical for the two firms, however, using numerical simulations shows that the downward sloping demand curves cause the better informed firm to charge a slightly lower price in equilibrium than the higher demand firm. The effect appears to be much smaller than the effect of the absolute level of variances through empirical observations.

As an illustration, suppose firm 1 receives a signal $S_1 \sim N(C, 1)$ and firm 2 $S_2 \sim N(C, 4)$, i.e firm 2 has much greater measurement error of the cost C . In the unit demand case, the Nash equilibrium is (5.16, 5.16) for both firms. In the case of a steep linear demand curve of $D(P) = 100 - 10P$, then the equilibriums change to (4.6, 5.6), see figure 3.3.




FIGURE 3.3: Numerically Derived Signal Reduction Best Response Curves - Firm 1 has $N(0,1)$ and Firm 2 has $N(0,4)$ signal error with a (-)ve sloping linear demand curve

This effect is understandable, albeit quite small relative to the absolute levels of the variances - if a firm is relatively more confident in their cost estimation, they benefit from making smaller signal reductions and capturing a greater share of the market.

3.5 OTC Market Bid-Ask Spreads

As detailed in Chapter 2, the true security value in a duopoly OTC market is analogous to the costs in a traditional firm/customer Bertrand competition. In this case, firms (the dealers) make a selling price of true value estimate plus a spread and additionally provide a buying price of their true value estimate minus a spread. Similar to above, each dealer attempts to estimate the true value of the security, which we model as receiving an independent sample V_i from a normal distribution. We assume the risk-neutrality of the market makers, so maximizing utility is the same as maximizing profits. Choosing the bid-ask spread is a symmetrical problem in which the firm makes a buy and a sell price for the asset of $(V_i - \delta_i, V_i + \delta_i)$ and they seek to maximize the expected payoff wrt to the value of δ_i .

This model of an OTC duopoly is set up as follows; two identical market-making firms (dealers) who compete in price for a homogeneous security are connected with a customer. The security has the same common value to both firms, which is unknown to both firms ex-ante. We also simplify the structure by assuming a unit demand from the customer, which is a more reasonable assumption in these markets - dealers often quote prices for a fixed quantity of the security known as 'Exchange Market Size' or EMS⁴ although the results from a linear demand curve are less tractable but numerical simulations produce similar effects. As in standard models, the dealers provide a firm buy and sell price to their connected clients. The dealers do not see the other prices before trading. The firm with the lowest price matches the customer buy order, and the highest buy price matches the customer sell order, and trades occur. After trades are matched, the true value is revealed and payoffs are realized. The payoff to the firms is price - cost, where price is the trading price and the cost can be viewed as the future true value (or future mark to market value, depending on the firm's time frame, which we assume is identical) of the security.

In exactly the same reasoning as above for the Bertrand duopoly problem, the market makers, who receive a signal (estimate) $V_i \sim N(V, \sigma)$ as to the true value V

⁴The London Stock exchange, for instance classifies every security with a EMS number, defined as the number of shares used to calculate the minimum quote size for each security and in fact 'must display bid and offer orders at the same time in at least Exchange market size and observe the maximum spread thresholds set out in parameters' - Stock Exchange Rule Book 2019

of the security, in equilibrium make a semi bid-ask spread of $\delta^*(\sigma) = \sigma\sqrt{\pi}$ in the symmetric information case.

This result is in line with the vast literature concluding that improved information about the true value reduces bid-ask spreads. The dealers also make positive profits in equilibrium, in contrast to the full informational model of Bertrand. This non-zero payoff in equilibrium questions the ubiquitous zero profit condition prevalent in many market microstructure models when considering multi-dealer competition.

As the dynamics are analogous to the simple 2-firm Bertrand case, the asymmetric case is also analogous. If two market makers have a different estimation ability to correctly value the security, they receive signals $S_1 \sim N(V, \sigma_1)$ and $S_2 \sim N(V, \sigma_2)$, then in equilibrium, both will make a price of $P(S_i) = S_i + \delta^*$ with $\delta^* = \sqrt{\frac{\pi}{2}(\sigma_1^2 + \sigma_2^2)}$, that is, both firms will use the same bid ask spread around their true value estimate.

This is even stronger evidence that public information can reduce spreads in the market as it has the greatest effect on the worst-informed firms, which affects the ability of the best-informed firms to exploit their informational advantage.

3.6 Client Reservation Prices

A central thread of this thesis is examining how different information sets between various agents change the equilibrium and payoffs in the system. The addition of customer reservation prices allows us to examine the asymmetric information case between customers and firms or clients and dealers in the OTC market case. Client-dealer information asymmetry is particularly relevant to the questions of transparency of information that are an important concern of policymakers. We consider boundedly rational clients in the sense that they attempt to estimate the true value but naively use this estimate as their reservation price.

Consider the standard Bertrand financial market competition model as above, with two identical market-making firms, homogenous security, and error signals $S_1 \sim N(V, \sigma_1)$ and $S_2 \sim N(V, \sigma_2)$ with a true common value V . We now include that the clients also set reservation prices that are based on their value estimate. The clients also do not have an exact knowledge of the value V and also attempt to measure it and are modeled by them receiving a signal $V_3 \sim N(V, \sigma_3)$, where they set their

reservation price. i.e., they set their reservation price at the naive level of their estimation of true value.

From the dealer's perspective, this formulation of the reservation prices is equivalent to viewing the clients as having a private asset valuation. These private valuations are normally distributed around the true value V .

The firm's spread problem now becomes:

Each firm $i=1,2$ makes a selling price (after receiving their signal V_i) of $P(V_i) = V + \varepsilon_i + \delta_i$ and has a payoff of $P(V_i) - V = (\varepsilon_i + \delta_i)$ if $(\varepsilon_i + \delta_i)$ is less than the $(\varepsilon_j + \delta_j)$ of the other firm and also less than the client reservation price ε_3 , zero otherwise. Formally we are looking for the probability of firm i , to win a competitive auction with one other firm.

$$P(p_i = V + \varepsilon_i - \delta_i < Y)$$

where $Y = \min\{\varepsilon_1 + \delta_1, \varepsilon_2 + \delta_2, \varepsilon_3\}$

Using the same methodology as in section 3.4.3, we can calculate the expected payoff to firm i , $i=1,2$ as:

$$E[\Pi_i] = E[\varepsilon_i + \delta_i] = \int_{-\infty}^{\infty} \frac{s}{\sigma_i} \phi\left(\frac{s - \delta_i}{\sigma_i}\right) (1 - \Phi\left(\frac{s - \delta_j}{\sigma_j}\right)) (1 - \Phi\left(\frac{s}{\sigma_3}\right)) ds \quad (3.25)$$

Unfortunately, this payoff function (3.25) does not have a nice algebraic representation, therefore, only numerical evaluation of the integral is possible. The dealers' objective is to maximize the payoff function (3.25) and the best response curves can be calculated numerically. For every choice of δ_j by firm j , firm i has a best response δ_i and is found by solving the first order condition of the payoff functions.

As an illustration, figure 3.4 is a plot of the 2 payoff surfaces with respect to the 2 delta choices with $N(0,1)$ signal errors and with customer reservation prices set at the naive valuation level. The best response curves are added which graphically suggest a single fixed point or point of mutual best responses and hence a Nash equilibrium.

The results of numerical simulations suggest identical effects to the non-reservation price duopoly case but with much lower payoffs in equilibrium. The reservation prices also cause spreads to narrow but is dependent on the level of information asymmetry




FIGURE 3.4: Payoff surfaces for the 2 firms with the best responses shown in red and blue

between firms and customers since adding a reservation price is similar algebraically to adding a third competitive firm into the competition.

3.7 Other Informational Asymmetries

Using numerical analysis, we can determine how prices respond to changes in customer information sets and how asymmetric dealer information affects the equilibrium when clients also measure the true value. We model the extent of an agent's information set by how accurately they can measure the unknown value by a single variable σ . We examine numerically how the equilibrium level spread changes for changes in relative customer/dealer informedness, that is, changes in the ratio of σ_M and σ_T

3.7.1 Client-Dealer Asymmetry in Monopoly

Consider the standard Bertrand financial market competition model as above, with only one monopoly dealer, homogenous security, and error signals $S_M \sim N(V, \sigma_M)$ with a true common value V . The clients receive a signal $S_T \sim N(V, \sigma_T)$, and again set their reservation price at the naive level of their estimation of true value.

The signal variance differences between the client and dealer lead to different prices and bid-ask spreads in equilibrium. In the limit, when the client is completely uninformed, and behaves as a liquidity trader, the dealer will price as large a spread as is possible.

The expected payoff to the monopoly dealer can be calculated as:

$$E[\Pi_D(\delta, \sigma_M, \sigma_T)] = \delta \Phi\left(\frac{-\delta}{\theta}\right) - \frac{\sigma_M^2}{\theta} \phi\left(\frac{\delta}{\theta}\right) \quad (3.26)$$

where $\theta = \sqrt{\sigma_T^2 + \sigma_M^2}$. The monopoly dealer's objective is to maximize equation 3.26. In the extreme case of the dealer being perfectly informed, $\sigma_M = 0$, and equation 3.26 is solved with an infinite spread as in the standard monopoly Bertrand case.




FIGURE 3.5: Monopoly dealer equilibrium spread levels varying with client relative signal variance

In figure 3.5, we fix the dealer's signal error to 1 and vary the client's signal standard deviation. We notice effects that reproduce results from the standard microstructure models. Clients better informed than the dealers correspond to the informed trader case in the Glosten and Milgrom (1985) model and dealers use increasingly higher spreads to protect themselves from their superior information. Similarly, as clients become worse informed, spreads increase as the dealer exploits their informational advantage. Spreads are lowest when clients are slightly worse informed than the dealer. The winner's curse effect - if the chosen spread is too low then there are negative payoffs, are compensated for by using wider spreads.

3.7.2 Client-Dealer Asymmetry in Duopoly

Consider a market setup as before with two dealers and one client, who sets a reservation price at their naive estimate of true value. In the case where both dealers are equally well informed, both dealers make the same spread, but the spread level is

now affected by the client's strategic actions, which are determined by their signal variance. Fixing the σ of the dealers at $\sigma = 1$ and varying the σ of the client produces the Nash equilibrium spreads shown in figure 3.6.




FIGURE 3.6: Dealer equilibrium spread level with client signal variance

The shape of the dealer's equilibrium spread levels are similar to the monopoly case but dramatically compressed. The equilibrium spreads are tighter when the client is semi-strategic and worse-informed than the dealers by choosing a reservation price over the non-reservation price levels. When clients are better informed, dealers respond by raising their spreads to protect themselves from the better-informed client.

3.7.3 Price Dispersion

Many papers look at price dispersion, and we can examine this in the duopoly case by calculating the variance of traded prices.

Proposition 3. *Given the duopoly OTC market with unknown common values described above and dealers receive independent signals $V_1 \sim N(V, \sigma_1)$ and $V_2 \sim N(V, \sigma_2)$ regarding the true value, then in equilibrium, where the 2 dealers use a symmetric pricing strategy: $Buy(V_i) = V_i - \delta_i^*$ and $Sell(V_i) = V_i + \delta_i^*$, then for any σ_1, σ_2 , the variance of the market buy and sell traded prices in equilibrium is $\theta^2 \left(\frac{\pi+1}{2\pi} \right)$*

Proof. Suppose $X_1 \sim N(\mu_1, \sigma_1)$ and $X_2 \sim N(\mu_2, \sigma_2)$. The first two moments of $Y = \min(X_1, X_2)$ are

$$E[Y] = \mu_1 \Phi\left(\frac{\mu_2 - \mu_1}{\theta}\right) + \mu_2 \Phi\left(\frac{\mu_1 - \mu_2}{\theta}\right) - \theta \phi\left(\frac{\mu_2 - \mu_1}{\theta}\right) \quad (3.27)$$

and

$$E[Y^2] = (\sigma_1^2 + \mu_1^2)\Phi\left(\frac{\mu_2 - \mu_1}{\theta}\right) + (\sigma_2^2 + \mu_2^2)\Phi\left(\frac{\mu_1 - \mu_2}{\theta}\right) - (\mu_1 + \mu_2)\theta\phi\left(\frac{\mu_2 - \mu_1}{\theta}\right) \quad (3.28)$$

where $\theta = \sqrt{\sigma_1^2 + \sigma_2^2}$, originally calculated by [Nadarajah and Kotz \(2008\)](#).

We examine the variance in the dealer traded selling prices (equivalently the client buy prices). As above, dealers make a selling price of $P_i = V + \epsilon_i + \delta_i$, where $\epsilon_i \sim N(0, \sigma_i)$, therefore $P_i \sim N(V + \delta_i, \sigma_i)$, $i = 1, 2$ and let $P = \min(P_1, P_2)$. Therefore, from the clients perspective, P is the trading price and the variance of P is :

$$\text{Var}[P] = E[P^2] - E[P]^2 \quad (3.29)$$

In equilibrium, $\delta^* = \theta\sqrt{2\pi}$ and $\delta_1 = \delta_2$ (by Proposition 1 and 2) and so the moments becomes:

$$E[P]^2 = \left(V + \theta\sqrt{2\pi} - \frac{\theta}{\sqrt{2\pi}}\right)^2 \quad (3.30)$$

$$E[P^2] = \left(V + \theta\sqrt{2\pi}\right)^2 + \frac{\theta^2}{2} - \left(\sqrt{\frac{2}{\pi}}V\theta + 2\theta^2\right) \quad (3.31)$$

$$\text{Var}[P] = \left(V + \theta\sqrt{2\pi}\right)^2 + \frac{\theta^2}{2} - \left(\sqrt{\frac{2}{\pi}}V\theta + 2\theta^2\right) - \left(V + \theta\sqrt{2\pi} - \frac{\theta}{\sqrt{2\pi}}\right)^2 \quad (3.32)$$

After some manipulation, the variance of P reduces to :

$$\text{Var}[P] = \theta^2 \left(\frac{\pi + 1}{2\pi}\right) \quad (3.33)$$

□

As the signal variance of either dealer goes up, the variance of the traded prices increases quadratically with theta, or linearly with the sum of the dealer signal variances, and in equilibrium depends only on the signal variance. The variance of the buying prices is equivalent by a symmetry argument.

3.8 Conclusion

We have looked at a particular case of the very well-studied Bertrand competition model with the added condition of an unknown common marginal cost and asymmetric information regarding the true cost. The common value case is particularly relevant for examining price formation in OTC financial markets and how the bid-ask spread is equivalent to the optimal bidding strategy in a first-price auction. We use the auction equilibrium results of [Wilson \(1969\)](#) to demonstrate how the bid-ask spread is affected by the clients' and dealers' information sets and extended our analysis to the asymmetric information case. This chapter's main narrative is that the extent of firms' information sets significantly affects their equilibrium pricing behavior. We examined the nature of the common value equilibrium with normal error signals and applied it to a competition duopoly. We found that firms with asymmetric information make the same level of signal reduction with a unit or linearly sloping market demand and how, in equilibrium, even worse informed firms can compete with better-informed firms. This result also explains why firms with worse informational costs are not driven out of the market. As with other papers on the subject, in particular, [Spulber \(1995\)](#), we find that informational uncertainty in the common value case is enough to resolve the Bertrand paradox of firms pricing at marginal cost and earning zero profits in expectation. We have assumed that competing dealers connect to identical clients, so they have the same form of the payoff function and use standard auction results. When clients are not homogenous, their network connections will affect the equilibrium payoffs and dealer strategies. These network effects offer an explanation for observed empirical facts that are otherwise difficult to explain, and we explore these effects in subsequent chapters.

Chapter 4

OTC Network Model

4.1 Introduction

In addition to the trading protocols detailed in Chapter 2, we now add a network component to the game. Before any trading occurs, each client must first select a subset of dealers and form potential trading relationships. This process forms the bilateral network links, and we develop a simple preferential attachment model that mimics this process. Preferential selection is the mechanism whereby clients express their preferences over the dealer set by choosing specific nodes over other nodes. This process of client link choices in the US corporate bond OTC market was highlighted by [Hendershott et al. \(2020\)](#), who examined the persistence of client links and the relationship distribution amongst clients and theorized that it was driven by linking costs and future trading expectations. The empirical result was that clients form a small number of persistent links to dealers, with some central dealers who execute a large proportion of the trade. A similar result was found by [Mallaburn et al. \(2019\)](#) in the UK corporate bond market.

The bipartite market model has been employed in a variety of contexts of economic exchange on networks, with and without intermediaries, see the review of intermediation networks, [Conderalli and Galiotti \(2015\)](#) for a brief summary of Intermediation networks and [Manea \(2009\)](#) for a comprehensive summary of the literature on bilateral trading on networks.

We take the market structure (who trades with whom) as given and, in this sense, are closer to the strand of literature that views agents as interacting on a fixed network, such as the seminal work by [Kranton and Minehart \(2001\)](#) (which shares some

ideas with this thesis using an ascending auction with private values) and others that were reviewed in Chapter 1. Gale and Kariv 2007; Blume et al. 2009; Gofman 2011; Manea 2011; Nava 2015; Condorelli, Galeotti, and Renou 2017; Choi, Galeotti, and Goyal 2017; Babus and Kondor (2018); Elliott 2015; Rahi and Zigrand 2013; Bramouille, Kranton, and D'Amours (2014). Like the random network matching models, the fixed network literature largely views decentralization as a restriction on the efficiency of trade. However, although we use a fixed network for examining equilibrium pricing, we think it is a novel approach to develop a method of generating a continuum of bilateral market networks using two parameters of preferential selection and client degree distribution to create networks with known degree centrality.

4.2 Basic Network and Graph Concepts

We start this chapter by summarizing some basic network concepts and definitions¹ that we use to construct the network OTC model. We also define some standard network formation models and briefly review the empirical OTC network literature. We then introduce a methodology to create a network with preferential selection using a fitness type model. The idea is to represent the trading connections between the clients and the dealers as a bipartite network, which we then use to define the dealers' payoff functions. This representation allows a much-simplified analysis of the equilibrium pricing levels.

4.2.1 Basic Definitions

The terms network and graph are typically used interchangeably in the literature; however, graphs are primarily used to refer to abstract mathematical objects, whereas, networks represent real-world objects in which the nodes represent entities of the system, and the edges represent the relationships between them. We start by defining a graph formally. Let us consider a finite set $V = \{v_1, \dots, v_n\}$ of unspecified elements and let $V \otimes V$ be the set of all ordered pairs $[v_i, v_j]$ of the elements of V . A relation on the set V is any subset $E \subseteq V \otimes V$.

¹The definitions in this section 4.2 are taken from Estrada and Knight (2015). Graph and Network Theory. In digital Encyclopedia of Applied Physics, Wiley-VCH Verlag GmbH Co. KGaA (Ed.). <https://doi.org/10.1002/3527600434.eap726>

We can define a simple graph as the pair $G = (V, E)$, where V is a finite set of nodes, vertices or points and E is a symmetric² and anti-reflexive³ relation on V , whose elements are known as the edges or links of the graph. A directed graph is defined as one where E is non-symmetric.

In many real world applications, the links or relationships between the nodes have weights and a more general graph definition is warranted. A weighted graph is defined as a quadruple $G = (V, E, W, f)$, where V is a finite set of nodes, $E \subseteq V \otimes V = \{e_1, \dots, e_m\}$ is a set of edges, $W = \{w_1, \dots, w_r\}$ is a set of weights such that $w_i \in \mathbb{R}$ and $f : E \rightarrow W$ is a surjective mapping that assigns a weight to each edge.

For a simple graph, an adjacency matrix A can be defined as $A_{ij} = 1$, if $e_{ij} \in E$, 0 otherwise. The degree of a node in a simple network is simply the number of links connected to it and in a directed network, we can define an in-degree and an out-degree representing the number of in-links and out-links that are incident on the node with a similar weighted adjacency matrix representation of the node edge relationship.

4.2.2 Types of Networks

The simplest type of graph or network is the tree. A tree of n nodes is a graph that is connected and has no cycles. The simplest tree is the path P_n . The path (also known as linear path or chain) is the tree of n nodes, $n - 2$ of which have degree 2 and two nodes have degree 1. We can find a spanning tree for any graph, a subgraph of this graph that includes every node and is a tree. A forest is a disconnected graph in which every connected component is a tree. A spanning forest is a subgraph of the graph that includes every node and is a forest.

An r -regular graph is a graph with $rn/2$ edges in which all nodes have degree r . A particular case of regular graph is the complete graph where every node is connected to each other. Another type of regular graph is the cycle, which is a regular graph of degree 2, i.e., a 2-regular graph, denoted by C_n . The complement of a graph G is the graph \overline{G} with the same set of nodes as G but two nodes in G are connected if and only if they are not connected in \overline{G} . An empty or trivial graph is a graph with no links. It is denoted as $\overline{K_n}$ as it is the complement of the complete graph.

²The relation E is symmetric if $[v_i, v_j] \in E \Rightarrow [v_j, v_i] \in E$

³The relation E is anti-reflexive if $[v_i, v_j] \in E \Rightarrow v_i \neq v_j$ and is reflexive if $\forall v \in V, [v, v] \in E$

A graph is bipartite if its nodes can be partitioned into two disjoint (non-empty) subsets $V_1 \subset V$, ($V_1 \neq \emptyset$) and $V_2 \subset V$, ($V_2 \neq \emptyset$) and $V_1 \cup V_2 = V$, such that each edge joins a node in V_1 and a node in V_2 . If all nodes in V_1 are connected to all nodes in V_2 , the graph is known as a complete bipartite graph and denoted by K_{n_1, n_2} , where $n_1 = |V_1|$ and $n_2 = |V_2|$ are the number of nodes in V_1 , V_2 respectively. Similar to the unimodal case, we can construct a matrix A , known as a biadjacency matrix, where the entry A_{ij} represents the link between nodes i and j .

Bipartite networks have proven to be a useful tool in many fields: condensed matter physics, socio-economic networks (firms and customers in traditional economic theory) and is the network structure that we examine in this thesis as the basis of the interconnectivity between clients and dealers in an OTC market.

4.2.3 Centrality Measures

Node centrality in a network is one of the many concepts that have been created in the analysis of social networks and then imported to the study of any kind of networked system. Measures of centrality can be either for each node – local centrality or for the whole network – global centrality. There is a distinction in centrality measures between directed and undirected networks. For directed networks, outgoing arcs are known as measures of influence, and incoming arcs are measures of support. In our client-dealer-directed network, the dealer links represent incoming price quote requests from the clients, and the outgoing links represent the trading quotes from the dealers to their connected clients.

Degree is the simplest of the node centrality measures that use the local structure around nodes and is the measure we use in identifying central dealers in the OTC network. In a directed network, degree is split into out-degree and in-degree, and this concept has been extended to weighted networks (Barrat et al. 2004, Newman, 2004) and labeled node strength. Other centrality measures, such as closeness centrality (the inverse sum of shortest distances) and betweenness centrality (the amount that a node lies on the shortest path between other nodes), have problems when applied in a bipartite network as many nodes are often disconnected, and it is the behavior of the competitive process with the dealer nodes that determine equilibrium pricing. Weighted degree centrality or node-strength then also naturally leads to how the

weights of an influencing node are distributed and finally to a global concentration measure.

The simple degree centrality of a node in a network was first considered as a centrality measure by Freeman (1979) in order to account for immediate effects taking place in a network. The degree centrality can be written as

$$k_i = \sum_{j=1}^n A_{ij} \quad (4.1)$$

In directed networks, we have in-degree and out-degree centrality of a node:

$$k_i^{in} = \sum_{j=1}^n A_{ij} \quad (4.2)$$

$$k_j^{out} = \sum_{i=1}^n A_{ij} \quad (4.3)$$

There are many other centrality measures that are useful depending on the context of the network studied. The closeness centrality, which measures how close a node is from the rest of the nodes in the network and in simple networks, the distance metric $d_{pq} = d(p, q)$ of 2 nodes, p and q is defined as the number of edges in the shortest path between them. Similarly for directed networks, there is an analogous directed distance measure which is a pseudo-metric as typically $d_{pg} \neq d_{qp}$, breaking the symmetry requirement for metrics. The closeness centrality, $CC(u)$, of a node u, is defined as $CC(u) = \frac{n-1}{s(u)}$, where the distance sum $s(u)$ is defined as $s(u) = \sum_{v \in G} d(u, v)$.

The betweenness centrality quantifies the importance of a node in relation to their position between other pairs of nodes in the network. It can be viewed as measuring the proportion of information that passes through the target node when there are communications between other pairs of nodes in the network. For every pair of nodes in a connected graph, there exists at least one shortest path between the nodes such that either the number of links that the path passes through (for unweighted graphs) or the sum of the weights of the links (for weighted graphs) is minimized. The betweenness centrality for each node is the number of these shortest paths that pass

through the node. It is defined as:

$$g(v) = \sum_{s \neq v \neq t} \frac{\sigma_{st}(v)}{\sigma_{st}} \quad (4.4)$$

Where $\sigma_{st}(v)$ is the total number of shortest paths from node s to node t and σ_{st} is the number of those paths that pass through v .

The Katz centrality is an important measure that can be viewed as an extension of the simple degree centrality where it seeks to include the influence of distant as well as close nodes. The Katz centrality Index can be defined, [Katz \(1953\)](#) as:

$$K_i = \{[(\mathbf{I} - \nu^{-1}\mathbf{A})^{-1} - \mathbf{I}]\mathbf{1}\} \quad (4.5)$$

where \mathbf{I} is the identity matrix, \mathbf{A} is the adjacency matrix, $\mathbf{1}$ is a column vector of 1's and $\nu \neq \lambda_1$ is an attenuation factor (λ_1 is the principal eigenvalue of \mathbf{A}).

Another type of centrality that captures the influence not only of nearest neighbors but also of more distant nodes in a network is the eigenvector centrality. This index was introduced by [Bonacich \(1987\)](#) and is the i^{th} entry of the principal eigenvector of the adjacency matrix.

Many of these centrality measures have severe interpretation issues when applied to a bipartite network as nodes in each disjoint subset do not connect to each other, hence we restrict the centrality measure to weighted and unweighted degree centrality. Notably, in most empirical research on centrality measures, the level of correlation between the different measures in real market structures is very high, see [Valente et al. \(2008\)](#).

4.2.4 Random Networks

The simplest and earliest model of a random graph was introduced by [Erdős and Rényi \(1959\)](#). A random graph in the Erdős Rényi model starts by considering some isolated nodes. Then, with probability $p > 0$ an edge is created between a pair of nodes. Consequently, the graph is determined only by the number of nodes, n and edges, m such that it can be written as $G(n, m)$ or $G(n, p)$.

4.2.5 Core-Periphery Networks

The intuitive concept of the set of nodes in a network having one subset that is more connected and central, and a peripheral subset of nodes that are less connected and less central has been used in economics since the 1950s. However, the definition of a core-periphery structure can and has been defined in numerous ways. The standard discrete model definition of a core-periphery network is a partition of the nodes into two subsets. One class of nodes forms a cohesive core sub-graph in which the nodes are highly interconnected, and the second class of nodes consisting of peripheral nodes that are loosely connected to this core, [Borgatti and Everett \(2000\)](#).

A limitation of this partition-based approach is the excessive simplicity of defining just two node types: core and periphery. This partitioning could be expanded to a three-class partition consisting of core, semi-periphery, and periphery, as world-system theorists have done. 'World-system refers to the inter-regional and transnational division of labor, which divides the world into core countries, semi-periphery countries, and the periphery countries.' Barfield, Thomas (1998).

In empirical studies of OTC markets, the network structure is found to have three main features. Firstly, a highly connected core of dealers, secondly, a small subset of dealers that intermediate a disproportionate amount of trade, and thirdly, client nodes with relatively few but persistent links to the dealers. These empirical observations of the OTC network topology (detailed in section [4.3](#)) suggest a type of core-periphery network structure by viewing the network as unimodal (core-periphery structures are not traditionally defined for bipartite networks).

4.2.6 Degree Distributions

The statistical distribution of the node degrees is a network characteristic that has received considerable attention in the literature. Let $p(k) = n(k)/n$ where $n(k)$ is the number of nodes having degree k in the network of size n . That is, $p(k)$ represents the probability that a node selected uniformly at random has degree k . The distribution of $p(k)$ versus k represents the degree distribution for the network. Depending on the generating mechanism of the network, the degree distribution could take many forms. A typical distribution which is expected for a random network of the type of

Erdős-Rényi is the Poisson distribution. However, a notably common characteristic of socio-economic networks is that many of them display some kind of ‘fat-tailed’ degree distributions. In these distributions a few nodes appear with very large degree while most of the nodes have relatively small degrees.

A typical example in the literature of these distributions is the power-law distribution, which is illustrated in Fig 4.1 taken from Estrada and Knight (2015), along with the Poisson distribution of an ER network. Other distributions such as lognormal, Burr, logGamma, Pareto, etc. (Foss et al., 2011) fall in the same category.

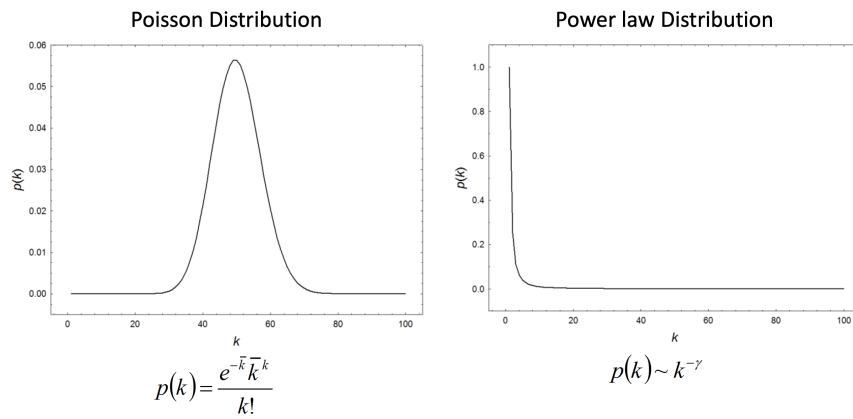


FIGURE 4.1: Degree distribution of Poisson and Power Law Networks

In power law networks, the probability of finding a high-degree node is relatively small in comparison with the high probability of finding low-degree nodes. These networks are usually referred to as ‘scalefree’ networks. The term scaling describes the existence of a power-law relationship between the probability and the node degree: $p(k) = Ak^{-\gamma}$ with γ referred to as a scale parameter.

Among the many possible degree distributions existing for a given network, the ‘scalefree’ one is one of the most ubiquitously found. Consequently, it is important to study a model that can produce random networks with such a degree distribution. That is, a model in which the probability of finding a node with degree decreases as a power-law of its degree. The most popular and one of the original models of these networks is the one introduced by Albert and Barabási (2002). Power-law degree distributions (with exponential tails) have been found to model the link structure in numerous OTC networks and other socio-economic network structures; however, ”knowledge of whether or not a distribution is heavy-tailed is far more important than

whether it can be fit using a power law”, [Stumpf and Porter \(2012\)](#). This observation is especially true in OTC competition networks.

4.3 Empirically Observed Market Networks

Most OTC markets, such as bonds, swaps, interbank lending, foreign exchange, real estate, and domestic energy, exhibit a stable core-periphery network structure. [Abad et al. \(2016\)](#), European interest rate, foreign exchange and credit derivatives and [Hendershott et al. \(2020\)](#), US corporate bonds, [Di Maggio et al. \(2017\)](#) and [Li and Schürhoff \(2019\)](#) in US municipal and corporate bonds, among others, find using TRACE and DTCC datasets, a power law link structure in numerous OTC financial markets. Although core-periphery structures are technically only defined for one node (unimodal) networks, the critical point is that core dealers tend to have higher centrality than periphery dealers. Abad et al. find that in financial markets⁴, roughly the same 10-15 dealers form the core, with the vast majority of trades having one of these core dealers as a counterparty. The largest 16 derivatives dealers, for instance, known as the ‘G16’ intermediate 53% of the total notional amount of interest rate swaps and 62% of credit default swaps. [Craig and von Peter \(2014\)](#) find a core-periphery structure in interbank lending markets.

There is an extensive literature of empirical observations of core-periphery networks and an equally large number of theoretical papers that attempt to explain the endogenous formation of a core-periphery structure. They range from some dealers having an ex-ante special advantage (for example, [Chang and Zhang \(2016\)](#)) to [Farboodi \(2015\)](#), who argues that the core-periphery structure is endogenously generated by counterparty default risk management. Notwithstanding, this configuration seems to be a defining feature of OTC bilateral markets. We find that the ubiquitous nature of this network structure is a key determinant of the more centrally located dealers’ spread behavior.

⁴These statistics were computed by [Abad et al. \(2016\)](#) using EMIR dataset

4.4 Preferential Attachment Network Models

Real-world network structures do not resemble the random network structures of the Erdos Renyi model, and core-periphery structures and power-law degree distributions are observed in many socio-economic structures, [Kolotilin and Panchenko \(2018\)](#). A common explanation for this is a random network formation process with preferential attachment. [Albert and Barabási \(2002\)](#) built the first theoretical model of network formation with a preferential attachment that differed from random network formation in two key ways. Firstly, growth - the random network model assumes a fixed number of nodes, whereas, in real networks, the number of nodes continually grows due to the addition of new nodes. Secondly, these new nodes prefer to connect to more popular nodes (as opposed to random connections) in a process known as preferential selection. These two characteristics, growth, and preferential attachment, define a network structure that more accurately models observed socio-economic networks and have dramatic implications for degree distributions. In OTC markets, these degree distributions determine the level of competition faced by the dealers and determine the amount of market power that firms have when pricing the product in a network and critical in affecting the equilibrium pricing behavior.

Although there is extensive literature on preferential growth networks, such as [Cooper and Frieze \(2003\)](#), who detail a more general model, and [Jackson and Rogers \(2007\)](#) with a hybrid homophily (nodes attach to similar nodes) preferential attachment model, the main mechanism of preferential selection (certain nodes are preferred in some way and results in a core-periphery structure in a unimodal network and concentrations of centrality in bimodal networks) is the basic generating feature of these economic market network models.

Although the network generating process is outlined above, we consider the network as fixed when we examine the pricing and equilibriums in a one-period game.

4.4.1 Bipartite Model with Preferential Selection

The traditional BA model of network formation in a unimodal network assigns a greater preference of newly created nodes to connect to nodes that already have a greater number of links. The ideas of the BA model of network formation can be

illustrated by applying them to our bilateral OTC network as follows: We start with m_0 dealer nodes and zero client nodes at $t = 0$. We add a new client node that connects to a dealer node with one link at each timestep. The probability that a link from this client node connects to dealer node D_i depends on the $k_i = \text{Deg}(D_i)$ as

$$P(k_i) = \frac{1}{m_0} + \left(\frac{m_0 - 1}{m_0} \right) \frac{k_i}{\sum k_i} \quad (4.6)$$

That is, the probability is skewed from pure randomness ($\frac{1}{m_0}$) by the degree of the dealer nodes $k_i = \text{Deg}(D_i)$ - the more clients that are connected to a dealer, the more likely they are to get selected by new nodes. As new clients enter the market, they increasingly make a connection to the higher degree dealer nodes. This process leads to a core-periphery structure for the client-dealer network, resulting in a small number of dealers having more links to the clients and executing a disproportionate amount of trade in the network - a result consistent with empirical evidence detailed in 4.3. After a large number of new client nodes have been formed, the resulting network has an almost identical structure to the BA unimodal model of network formation with a degree distribution of the dealer nodes that have a power-law distribution. The BA model has been subject to some criticism with regard to real-world networks. BA themselves recognize that the initial configuration is critically important to the eventual link distribution of the network. This model makes it difficult for any new nodes to become significant once the network is established (it is a pure rich-get-richer model). In a competitive market network with firms and clients, this feature is not desirable as new entrant firms can often become popular (aggressive pricing, superior technology, marketing, etc.).

Although existing degree (popularity) is the source of preferential selection amongst new client nodes in the BA model, we examine a more general situation where there is some attribute of the dealer nodes that allows the clients to rank them in terms of preference. In terms of modelling, this does not change the analysis. Suppose we have m_0 dealer nodes that a new client node could connect to. If there was no preference, then each dealer node gets randomly selected with probability $\frac{1}{m_0}$. We add preferential selection by allowing some nodes to have a greater probability of being selected.

Suppose dealer node 1 is preferred by clients in some way (perhaps because of size, risk management or marketing strategy or some combination of defining features) and this is modelled by dealer 1 having an excess probability of being chosen of α , where $0 \leq \alpha \leq \frac{m_0-1}{m_0}$. We assume that α is constant. The preferred dealer node 1 has a probability $p(D_1)$ of being selected by a new client node as:

$$p(D_1) = \frac{1}{m_0} + \alpha \quad (4.7)$$

The remaining $m_0 - 1$ dealer nodes now have a probability of being selected as:

$$p(D_i) = \frac{1}{m_0} - \frac{\alpha}{m_0 - 1} \quad (4.8)$$

This formulation is more flexible than just using node degree for preferential selection. We can control the exact amount of preferential selection in the network, which allows us to examine the relationship between this preferential selection parameter α , (which also maps linearly onto weighted degree centrality), and the equilibrium prices charged in the network.

More generally, if we assign a constant probability p_i to each dealer node being selected, then after the creation of N client nodes, the expected degree distributions follows: $E[\deg(D_i)] = Np_i$. and the ratio of the expected degree of any two nodes, D_i and D_j is: $\frac{E[\deg(D_i)]}{E[\deg(D_j)]} = \frac{p_i}{p_j}$

We formulate these ideas of network formation and preferential selection in more detail in section 4.8.

4.4.2 Higher Degree Client Node Networks

So far, we have looked at each new client selecting only one link, but the clients may select more links than this in order to receive a more competitive price. The amount of connectivity (degree) of clients and their degree distribution may affect the equilibrium pricing level. The price quoted in a 1-link monopoly structure is often different from that quoted in a multi-link connection with competing firms. Indeed, dealers' price quotes depend on the amount of competition they face with other dealers and, hence, the quotes charged in the resulting equilibrium strategies.

This simple model can also incorporate this client characteristic by allowing a distribution over the clients' link structure. Suppose clients have a degree that is distributed with a probability mass function f , such that $P(k_i = s) = f(s)$ where k_j ranges from $1 : m_0$, that is, there is a probability mass function that determines if a client connects to only 1 dealer (monopoly) up to m_0 dealers (fully competitive). Suppose a new client node is created that has a degree of k links and connects to the dealer set D , $|D| = m_0 \geq k$, we want to calculate the expected degree of the preferred and non-preferred nodes in D . The probability that a link from this client node connects to a preferred dealer node D_1 is calculated and then used to calculate the expected degree. The probability p_i of being selected by the i^{th} link (and not having been selected by the previous $(i - 1)^{th}$ links is defined recursively as:

$$p_0 = 0$$

$$p_i = \prod_{j=1}^{i-1} (1 - p_j) \left(\frac{1}{(m_0 - j + 1)} + \alpha \right)$$

And so the probability $P(k)$ of the preferred dealer node receiving a link from a client with degree k is:

$$p(k) = \sum_{i=1}^k p_i = \sum_{i=1}^k \prod_{j=1}^{i-1} (1 - p_j) \left(\frac{1}{(m_0 - j + 1)} + \alpha \right) \quad (4.9)$$

Suppose that N client nodes have been created (with fixed m_0 dealer nodes), and let $\beta_1, \dots, \beta_{m_0}$ be the proportions of the N clients that have a degree of $1, \dots, m_0$, where $0 \leq \beta_i \leq 1$ and $\sum \beta_i = 1$, then the expected degree $k_1 = Deg(D_1)$ of the preferred dealer node is:

$$E[k_1] = N (\beta_1 p(1) + \dots + \beta_{m_0} p(m_0)) = N \sum_{i=1}^{m_0} \beta_i p(i) \quad (4.10)$$

Various distributions of $\beta = [\beta_1, \dots, \beta_{m_0}]$, of the client degree, give rise to various familiar economic network structures. If we classify each generated network by parameters (α, β) , then the particular network $(\alpha = 0, \beta = \{1, 0, 0, \dots, 0\})$, case, corresponds to a bilateral star network where each client connects to only one dealer randomly (evenly

to the dealers in expectation). This is a monopoly market with multiple monopoly dealers, each with their own client base. At $\alpha = \frac{m_0-1}{m_0}$, (it's maximum value), and $\beta = \{0, 0, \dots, 0, 1\}$ corresponds to the complete network, the totally competitive market. By varying these parameters of client degree β and preferential selection α we can generate the entire space of possible competitive networks.

These two generating characteristics, the distribution of the client links to the dealers β and the amount of preferential selection α for the dealers by the clients when forming their links, constitute the basis of our OTC network model. We will also map these parameters onto a network centrality measure in section 4.8.

4.5 N Dealer, M Client Network Game

Following the network structure described in section 4.4.1, how the clients connect to dealers (in our model represented by the tuple (α, β)), is the key element in determining the amount of competition in each potential client-dealer transaction and hence the equilibrium pricing choices of the dealers. We introduce two types of agents: Dealers $D_i : i = 1, \dots, N$ Clients $C_j : j = 1, \dots, M$. Due to restraints of search costs, each client node C_i , maintains contact with a finite number $k_i = \text{Deg}(C_i) (\leq N)$ of dealer nodes from whom they request quotes. k_i is the number of links with the neighborhood set $N[C_i]$ that can be maintained given the clients' search and maintenance cost constraints. The intuition is that it is a lengthy (and hence costly) process to create and maintain trading relationships with dealers, limiting the number of connected dealers.

We characterize each client $C_i(V_i, N[C_i])$ with 2 parameters, V_i , which is the clients' value of the asset and $N[C_i]$, the neighbourhood set (sub-network) of dealers $\{D_1, \dots, D_{k_i}\}$ that the client is connected to, in order to request prices.

The links characterize potential trading relationships between the clients and the connected dealers, and the OTC game of the set of $\{D_i\}$ dealers and $\{C_i\}$ clients can be represented by a bipartite graph game $G = (D, C, E, \Pi)$, where E denotes the links or trading relationships between the dealers and clients. Cardinality $|C| = M$ $|D| = N$ and Π , the payoff functions for each dealer given their pricing choices p_i .

The biadjacency matrix of G is a $|C| \times |D|$ matrix E , where $e_{ij} = 1$ if a trading relationship exists and 0 if not. $|C| \gg |D|$ or $M \gg N$, (there are many more clients than dealers). We use the simplified notation that $G(C_j, N[C_j])$ is the competitive price game comprising client C_j and their connected dealer neighbourhood set $N[C_j]$. $\Pi_{D_i} G(C_j, N[C_j])$ is the payoff to D_i in this sub-network (with dealer prices $p = [p_1, \dots, p_n]$).

This chapter is primarily concerned with the network structure and connections between the clients and dealers. However, in order to simplify the OTC competition network, we introduce some game structure to the bilateral network. Clients connect to the dealers and form a network structure given by (α, β) . We assume that α and β are common knowledge.

The trading protocols follow the usual structure. The clients request a trading price from their connected dealers (their neighbourhood sets $N[C_i]$) and each dealer D_i quotes a trading price p_i to their neighbourhood client sets $N[D_i]$. The quoted trading price represents the dealers' strategy in this game. Each client and dealer has a valuation V_i for the homogenous good. The clients choose the best price from these dealers, and trade occurs in a Bertrand price competition. process.

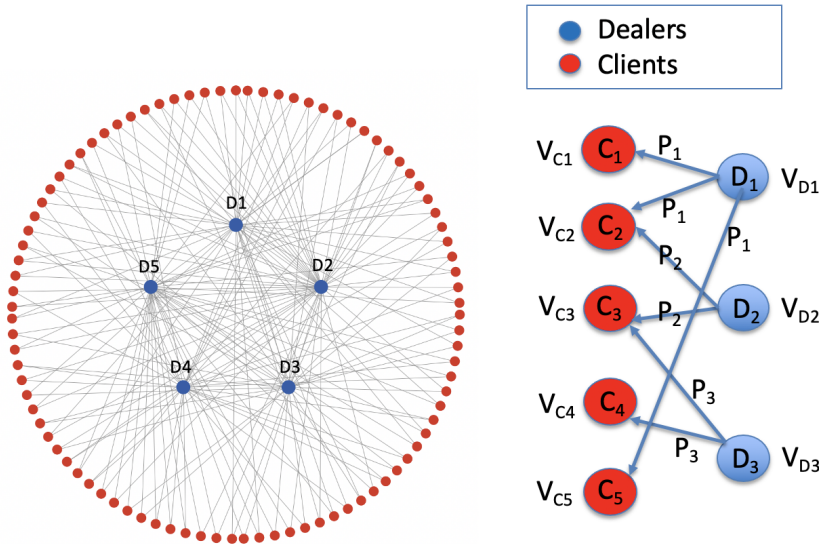


FIGURE 4.2: Clients C_1, \dots, C_5 and associated neighbourhood sets and valuations

Figure 4.2 shows a typical OTC market network. The second image illustrates

a subsection of the main network involving clients C_1, \dots, C_5 who are connected to dealers D_1, D_2, D_3 . The clients submit requests for trading prices to these connected dealers and these dealers respond with prices p_1, p_2, p_3 . Both clients and dealers have a valuation of the asset V_j

The payoff to D_1 is the payoff from interacting with his neighbourhood set $N[D_1] = \{C_1, C_2\}$ and hence is the sum of the payoffs of the sub-games $G(C_1, N[C_1])$ and $G(C_2, N[C_2])$. The probability of trading and the trading price is affected by the prices p_i of the dealers and the payoff to each dealer is a function of each dealer's valuation level which we explore in greater detail in the next chapter.

4.6 Network Payoffs

Let there be N dealers and M clients connected in a network represented by $\Gamma = [\{C, D\}, A]$ where C, D are the sets of client and dealer nodes, connected via a biadjacency matrix A . Let $G = [C_i, N[C_i], A^*]$ be a subnetwork consisting of a client C_i and its connected dealer set $N[C_i]$.

We introduce the notation that $E[\Pi_{D_i} G(C_j, D_1, \dots, D_m | \delta_1, \dots, \delta_m, \sigma_M, \sigma_T)]$ is the expected payoff to dealer D_i when connected to a client C_j , who is also connected to m dealers D_1, \dots, D_m and each of these dealers uses semi bid-ask spreads of $\delta_1, \dots, \delta_m$. The set of connected dealers of C_j is the neighbourhood set $N[C_j]$. It is understood that the expectation is conditional on the action set $\{\delta_1, \dots, \delta_m\}$ of the dealers and error variances σ_M and σ_T and so we will drop this from the notation for brevity. Generally, the clients are connected to different subsets of dealers and so the expected payoff of dealer i is the sum of payoffs over all the clients they are connected to and each interaction with each client may have different numbers of competing dealers. Let $C_i \in \{C_1, \dots, C_M\}$ and let $N[C_i]$ be the neighbourhood set of dealers of client C_i . The expected payoff or profit of dealer i in the entire network game G , $E[\Pi_{D_i}(\Gamma)]$ is:

$$E[\Pi_{D_i}(\Gamma)] = \sum_{s \in N[D_i]} E[\Pi_{D_i} G(C_s, N[C_s])] \quad (4.11)$$

where $N[D_i]$ is the neighbourhood set of dealer i , D_i . $N(C_s)$ is the neighbourhood of client C_s and each C_s is a member of $N(D_i)$. This expression is just the sum of

payoffs over all subgames with connected clients of dealer i .

For intuition, consider a subset of a dealer network that comprises $\{D_i, D_1, D_2, D_3, D_4\}$ with a topology described in figure 4.3.




FIGURE 4.3: Example of a 5-dealer market with clients C_1, C_2, C_3 and associated neighbourhood sets

It is clear that dealer D_i 's payoff depends on both its neighbourhood set $N[D_i]$ and the neighbourhood sets of this neighbourhood set, $N[N[D_i]] = \{N[C_1], N[C_2]\}$. Although D_i is not connected to C_3 , the pricing decisions of D_3 with C_3 affect the pricing decisions with D_3 and C_2 , which in turn affects the pricing decision of D_i with C_2 . Through this mechanism, pricing and information flows through the network and causes interconnectivity of pricing strategies.

Equation 4.11 is the sum of payoffs from each individual client that dealer D_i is connected to and each term of the sum has the form of equation (2.8). The dealers all seek to maximize this expected payoff function (4.11) in equilibrium by choosing an appropriate bid-ask spread δ_i that is used in all of their sub-network games.

This formulation is the essence of the OTC game structure - a sum of competition payoffs with multiple clients in a network configuration with each dealer using a single choice variable δ_i , the bid-ask spread.

4.7 Weighted Biadjacency Matrix Representation

The links that represent potential trading relationships between clients and dealers have been characterized as an unweighted bipartite network, where each link has a weight of 1, which represents a potential trading relationship. This unweighted

network can be converted into a weighted bipartite network, which reduces the total number of possible market topologies (only a limited number of ways to connect to the dealer set) and allows us to measure the effects of dealer centrality more easily. The interactions between the set of D dealers and C clients through the links can be represented by a network game $G = (D, C, E, \Pi)$, where E denotes the links or potential trading relationships between the dealers and clients. Cardinality $|C| = M$, $|D| = N$ and Π , the payoff functions. The biadjacency matrix of G is a $|C| \times |D| (= M \times N)$ matrix E , where $e_{ij} = 1$ if a trading relationship exists and 0 if not. $|C| \gg |D|$ or $M \gg N$.

We introduce a more compact form $G^* = (D, C^*, E^*)$, where the clients are now partitioned into sets with the same topological link structure. i.e., if two clients are both connected to the same set of dealers, we can use one node and a weighted link. This is allowable because we assume that the expected payoff to a dealer from two identical clients is the same. The biadjacency matrix of G^* is a weighted biadjacency matrix E^* , where $e_{ij} = \epsilon_{ij}$ are the proportion of clients with a certain link structure and $0 \leq \epsilon_{ij} \leq 1$.

Monopoly case

Consider a monopoly market network (star network) with M homogenous clients and a single dealer; we can represent this as a weighted network with 2 nodes: a client node and a dealer node as in figure 4.4.




FIGURE 4.4: Strategic equivalence of monopoly market, M clients

Node $C1$ represents the only strategic choice of the clients. The payoff for the dealer $D1$ only depends on the weight M . We can equivalently express this as $M = 1$, which is 100 % of clients connect to the dealer $D1$ with one link. Importantly, we assume the expected payoff from each client with the same link structure is the same,

then the expected payoff in the network game G from this node is the sum of the expected payoffs from each of the M identical clients, i.e:

$$E[\Pi_{D_i}(G)] = \sum_{j=1}^M E[\Pi_{D_i}(C_j)] = ME[\Pi_{D_i}(C1)] \quad (4.12)$$

Duopoly Case




FIGURE 4.5: Strategic equivalence of duopoly market, M total clients

Similarly with a 2-dealer market, here clients can either have 1 link to either dealer or connect with both in a competitive process as in figure 4.5.

The link weights are the proportions of M that have that particular topology. $\alpha + \beta + \gamma = 1$ and $C1$, $C2$, $C3$ represent all the possible strategic choices of the clients. i.e. $C1$ represents the clients' choice to have a single link to dealer 1, $D1$. $C2$ is the set of clients with a single link to $D2$ and $C3$ represents the set of clients that connect to both. If clients have no preference for a particular dealer, then $\alpha = \beta$ in expectation.

If clients choosing the same strategy (topology) are assumed to be homogeneous, then no information about the network is lost in this transformation.

General case

Every possible bipartite network adjacency matrix $E(N \times M)$, can be represented with a new single matrix $E(N \times 2^N - 1)$ with different biadjacency weight parameters representing the partition of the client topology.

In a general network of N dealers and M clients with each client connecting to k dealers, there are C_k^N possible nodes (+ n dealer nodes) in the network (where C_k^N

are the binomial weights). Each dealer is involved in $2^N - 1$ distinct auctions (some will be empty depending on the network)

Therefore the union of all these networks ($k=1:N$) represents the network of all strategic possibilities for the M clients - these are the only ways to connect to a dealer subset.

This union network U has total number of nodes or client strategic choices:

$$|U| = \sum_{i=1}^N \mathbf{C}_i^N = 2^N - 1 \quad (4.13)$$

4.7.1 Dealer Payoff in Weighted Network Representation

Again, using the notation that $G(C_j, N[C_j])$ is the competitive game played between the client C_j and its neighbourhood set of connected dealers $\{D_1, \dots, D_k\}$ and $E[\Pi_{D_i}(G)]$ is the expected payoff for dealer i in game G . Let the set of dealers be a vector $D = [D_1, \dots, D_n]$ and the set of possible clients' connective strategies $C = [C_1, \dots, C_p]$ where $p = 2^N - 1$, i.e. we group together all the clients that have the same topology or equivalently, identical neighbourhood sets. Let A be the biadjacency matrix of connections where a_{ij} is the weight of the connection $0 \leq a_{ij} \leq 1$ and represents the proportion of clients connected in this way. The expected profit of a dealer is the sum of expected profits from each of its connected client groups.

Let there be N dealers and M clients and let α_i be the proportion of clients connected to each possible dealer subset, of which there are $2^N - 1$ possible subsets (the binomial permutations). See section 4.7. The payoff to dealer i who is connected to multiple clients C_1, \dots, C_k , where each client C_j is connected to a subset of dealers $N[C_j]$ is:

$$E[\Pi_{D_i}(\Lambda)] = M \sum_{j=1}^{2^N-1} \alpha_j E[\Pi_{D_i} G(C_j, N[C_j])] \quad (4.14)$$

Which is the sum of payoffs from each individual client that dealer D_i is connected to and each term of the sum has the form of equation (2.8). The dealers all seek to maximize this expected payoff function (4.14) in equilibrium by choosing an appropriate bid ask spread δ_i that is used in all of their sub-network games.

In general and in matrix form, the expected payoff for the dealers:

$$E[\Pi_D(G)] = M \mathbf{A} \mathbf{E}[\Pi_D(\mathbf{C})] \quad (4.15)$$

, where $D = [D_1, D_2, \dots, D_N]^T$, $E[\Pi_D(C)] = [E[\Pi(C_1)], E[\Pi(C_2)], \dots, E[\Pi(C_p)]]^T$

M is the number of clients and \mathbf{A} =weighted biadjacency matrix of connections, $N \times (2^N - 1)$ and $E[\Pi(C_i)]$ is the expected profit from one of the possible client auction games C_i . If the dealer is not connected to a client, we set the payoff in this game to zero.

Example in a 3-dealer network

There are $2^3 - 1 = 7$ client strategic possibilities in this game (represented by the columns of the biadjacency matrix)

$$A = \begin{bmatrix} \alpha_1 & 0 & 0 & \beta_1 & \beta_2 & 0 & \delta_1 \\ 0 & \alpha_2 & 0 & \beta_1 & 0 & \beta_3 & \delta_1 \\ 0 & 0 & \alpha_3 & 0 & \beta_2 & \beta_3 & \delta_1 \end{bmatrix}$$

where the weights $\sum \alpha_i + \beta_i + \delta_i = 1$

$$C = [E[\Pi(G(D1))], E[\Pi(G(D2))], E[\Pi(G(D3))], E[\Pi(G(D1, D2))], \\ E[\Pi(G(D1, D3))], E[\Pi(G(D2, D3))], E[\Pi(G(D1, D2, D3))]]$$

Consider dealer D1, represented by the first row of matrix A. D1 is connected to C1, C4, C5 and C7. The expected payoff is a function :

$$E[\Pi_{D1}] = \alpha_1 E[\Pi_{D1}(C1)] + \beta_1 E[\Pi_{D1}(C4)] + \beta_2 E[\Pi_{D1}(C5)] + \delta_1 E[\Pi_{D1}(C7)] \quad (4.16)$$

and in terms of distinct games:

$$E[\Pi_{D1}] = \alpha_1 E[\Pi_{D1}(G(D1))] + \beta_1 E[\Pi_{D1}(G(D1, D2))] \\ + \beta_2 E[\Pi_{D1}(G(D1, D3))] + \delta_1 E[\Pi_{D1}(G(D1, D2, D3))]$$

Where for example $G(D1, D3)$ is the distinct auction game involving dealers 1 and 3 and the $E[\Pi_{D1}(G(D1, D3))]$ is the expected payoff for D1 in game $G(D1, D3)$

4.8 Simplified Preferential Attachment Network

4.8.1 1 and n Degree Client Case

To simplify client network choices but maintain the sensitivity between higher and lower client degree choices, we assign each client to either a low or high degree class. The low degree class has only one link to a dealer, and the high degree class connects to all of the dealers in the network. These simplifications model the fact that some clients only choose limited price competition due to link creation costs, whereas other clients choose many firms to quote with. Therefore, firms would like to have a different pricing strategy depending on how much competition they face in each client interaction. This simplification is consistent with the observed power-law link distributions observed in OTC markets, [Hendershott et al. \(2020\)](#) and [Mallaburn et al. \(2019\)](#).

The weighted biadjacency matrix has a simple form and becomes:

$$A = \begin{bmatrix} \alpha_1 & 0 & 0 & 0 & 0 & 0 & \gamma \\ 0 & \alpha_2 & 0 & 0 & 0 & 0 & \gamma \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \gamma \\ 0 & 0 & 0 & 0 & 0 & \alpha_n & \gamma \end{bmatrix}$$

which is an $n \times (n+1)$ matrix and $\sum \alpha_i + \gamma = 1$

The payoff vector for the N dealers with M clients is: $E[\Pi(\mathbf{D})] = M\mathbf{A}\mathbf{C}$, where $\mathbf{D} = [D_1, D_2, \dots, D_N]^T$ and $\mathbf{C} = [E[\Pi(C_1)], E[\Pi(C_2)], \dots, E[\Pi(C_{N+1})]]^T$

The payoff for a dealer i is:

$$E[\Pi_{Di}] = \alpha_i E[\Pi_{Di}(G(D_i))] + \gamma E[\Pi_{Di}(G(D_1, \dots, D_N))] \quad (4.17)$$

which is the expected payoff for dealer i from 2 distinct competitions - the monopoly game $G(D_i)$ and the complete game $G(D_1, \dots, D_N)$.

(In equilibrium, generally $E[\Pi_{Di}(G(D_1, \dots, D_K))] \neq E[\Pi_{Dj}(G(D_1, \dots, D_K))]$ - as the payoff and choices for each dealer in equilibrium are affected by the spread choices in

all the other subgame competitions that each of the dealers are involved in.)

We now also add preferential selection to the simplified matrix representation by allowing one of the dealers to be more popular. Each dealer with 1 link is equally likely to be selected, with probability $1/n$. Suppose dealer n , D_n is preferred and its' probability of being selected is increased by α . The probability p of D_n being selected is $p = \frac{1}{n} + \alpha$ and the other dealers $[D_1, D_2, \dots, D_{n-1}]$ have a probability of $\left(\frac{1}{n} - \frac{\alpha}{n-1}\right)$ of being selected. Suppose also that a proportion β of clients have degree 1 and $(1-\beta)$ have degree n . Our expected biadjacency matrix now looks like:

$$A = \begin{bmatrix} \beta \left[\frac{1}{n} - \frac{\alpha}{n-1} \right] & 0 & 0 & \cdot & 0 & 0 & (1-\beta) \\ 0 & \beta \left[\frac{1}{n} - \frac{\alpha}{n-1} \right] & 0 & \cdot & 0 & 0 & (1-\beta) \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & (1-\beta) \\ 0 & 0 & 0 & \cdot & 0 & \beta \left[\frac{1}{n} + \alpha \right] & (1-\beta) \end{bmatrix}$$

The biadjacency matrix has 2 parameters, α and β , where:

α characterizes the amount of preferential selection of the preferred dealer and maps linearly to weighted degree centrality, $0 \leq \alpha \leq (n-1)/n$ and determines the distribution of centrality in the network

β represents the proportion of low degree clients in the network and determines the average centrality in the network $0 \leq \beta \leq 1$

By varying α and β , we can examine how the equilibrium prices change for different types of competition market networks. $\alpha=0$ corresponds to no central node, $\alpha=1$ is 1 central dealer trading exclusively with all low degree clients and equally participating with the high degree clients (weighted degree centrality =1). $\beta=0$ corresponds to a complete network, $\beta=1$, corresponds to the monopoly network.

NB:

$(\alpha, \beta) = (0, 0)$ corresponds to the complete network (total competition)

$(\alpha, \beta) = (1, 1)$ corresponds to the star network (monopoly market)

The idea is to move between these 2 extreme networks by varying these parameters

We can now use this variable network representation in the payoff function of a dealer in this network auction game: The peripheral dealers $\{D_i\}$ have a payoff function:

$$E[\Pi_{D_i}] = \beta \left[\frac{1}{n} - \frac{\alpha}{n-1} \right] E[\Pi_{D_i}(G(D_i))] + (1 - \beta) E[\Pi_{D_i}(G(D_1, \dots, D_n))] \quad (4.18)$$

and the central dealer D_n has an expected payoff:

$$E[\Pi_{D_n}] = \beta \left[\frac{1}{n} + \alpha \right] E[\Pi_{D_n}(G(D_n))] + (1 - \beta) E[\Pi_{D_n}(G(D_1, \dots, D_n))] \quad (4.19)$$

where $E[\Pi_{D_i}(G)]$ is the expected payoff for D_i from the game G

We could equally have chosen 2 links and $n-1$ links as the 2 examples to simplify the network problem, but this is complicated by the connection permutations and the results are identical. The important features are captured in this representation, being the 2 extremes of the client game choices, which is a blend of monopoly and fully competitive payoffs.

In the following chapter, we specify the exact form of $E[\Pi_{D_i}(G)]$, which is the Bertrand price competition payoff with unknown common valuations.

The expected weighted degree centrality $DC(D_n)$ of the central dealer described above given the parameters (α, β) can be calculated as:

$$DC(D_n) = E[k_n] = \beta \left[\frac{1}{n} + \alpha \right] + (1 - \beta) \quad (4.20)$$

and the peripheral dealers, D_p ($p=1, \dots, n-1$), have an expected weighted degree centrality as:

$$DC(D_p) = E[k_p] = \beta \left[\frac{1}{n} - \frac{\alpha}{n-1} \right] + (1 - \beta) \quad (4.21)$$

Where as before n is the (fixed) number of dealers, and (α, β) are the network generating parameters as described above. For a fixed preferential selection α , the centrality of all dealer nodes is an increasing linear function of the amount of low degree clients β and similarly for a fixed β the centrality is an increasing linear function of α for the core dealer and a decreasing linear function for the periphery dealers.

Although both α and β affect the centrality calculation, they affect it in different ways. β affects the global level of dealer centrality and network density, whereas




FIGURE 4.6: Relationship between α and weighted degree centrality with 50% low and high degree clients




FIGURE 4.7: Relationship between α and relative degree centrality $DC(D_n)/DC(D_p)$ of central to peripheral dealers with various proportions of low degree clients

α affects how unequally the links and centrality are distributed. Figure 4.7 shows the centrality of the central dealer, normalized by the centrality of the non-preferred dealers. Our next chapter examines the centrality premium effect, that is how central and peripheral dealers make different prices in equilibrium that is related to their relative centrality in the network and so we will typically fix β and vary the amount of relative preferential selection (and hence the centrality measures of the central and peripheral dealers) α to better understand the origins of this centrality effect.

4.9 The OTC Network as a Static Bayesian Game

We can represent our OTC market network game model as a Static Bayesian game as in Chapter 2 with the addition of the network generating parameters. All other protocols remain the same. A Bayesian game is a game in which the players have

incomplete information of the other players' strategies or payoffs, but they have beliefs over other players' types with a known probability distribution. With a network setting, types now refer to both network position and signal realization. A Bayesian game can be converted into a game of complete but imperfect information under the common prior assumption. Each player knows the probability distribution over types, and this itself is common knowledge among all the players. Firstly, clients connect to dealers with a probability mass function $p(\cdot)$ that is determined by the generating parameters α, β . Secondly, dealers and clients receive an independent signal V_i as to the asset's true common value, and this pre-game process determines player types. Dealers quote a buy and a sell price to the clients with a bidding strategy that maps their signal (or estimation of the true value) onto real numbers, such that,

$$\begin{aligned} P^{BID}(S_i) &= S_i - \delta_i : S_i \rightarrow \mathbb{R} \\ P^{ASK}(S_i) &= S_i + \delta_i : S_i \rightarrow \mathbb{R} \end{aligned} \tag{4.22}$$

Where, $S_i \in (-\infty, \infty)$, $\delta_i \in [0, \infty)$.

We assume a symmetric bidding strategy, (each dealers' bidding function has the same linear functional form) and is common knowledge, as is the probability distributions of signals and client link distributions in the network and, as is usual with Bayesian games, we assume risk neutrality, so maximizing utility is the same as maximizing payoffs.

(i) Players: M client nodes and N dealer nodes

(ii) Action spaces: Dealers: $\delta_i \in \mathbb{R}^+$, Clients: $\{0, 1\}$ where 0 indicates no trade and 1 indicates a unit trade with a connected dealer depending on their reservation price

(iii) Distribution over player types:

Dealers: dealer i's type space is a 2-tuple $(S_i, N(D_i))$ where each signal S_i is IID normal RV drawn from $N(V, \sigma_M)$ and dealer i's neighbours (clients) are drawn from a known distribution.

Clients: Client j's type space is a 2-tuple $(S_j, N(C_j))$ where each S_i is IID normal RV drawn from $N(V, \sigma_T)$ and client j's neighbors $N(C_j)$ (dealers) are also drawn from a known distribution.

Proposition 4. *In the OTC network game described above, the dealers can form an expression for their expected payoff function without a full knowledge of the network connections and it is sufficient to know the generating parameters of the bipartite network (α, β) .*

Proof. Consider D_i , dealer i . Nature chooses their signal, S_i and their local network of connections to clients, $N[D_i]$. Dealer i knows their own signal (their estimate of true value) and their neighbourhood (the number of clients they are connected to). They don't know the signal realizations of the other dealers and clients and importantly, doesn't know the network properties of other dealers or clients. Importantly, they doesn't know the network properties of their connected clients. We use the common prior assumption, [Harsanyi \(1967\)](#), and dealer D_i then forms a belief about the type of each client and dealer in the market, conditional on their own signal and their neighbourhood set realization (signals of each client and dealers and beliefs about the likely network connections).

As before, let $\Pi_{D_i}(C_j, \{N(C_i), \delta_i, \delta_{-i}\})$ be the payoff function for dealer D_i , in an auction game with client C_j who has a local network of dealer connections $\{N(C_j)\}$. Dealer D_i now forms an expectation of this neighbourhood set, given the network link distribution, which is common knowledge. He also forms an expectation about the signals of other players.

In a bipartite network, each client has at most $2^N - 1$ possible connections to different dealer sets (identical neighbourhoods) and so must belong to 1 of these client types. Each of these $2^N - 1$ clients types occurs with probability $p(\cdot)$ which is determined by the probability mass function that we are taking as common knowledge.

The expected payoff to dealer i in the whole network Λ , using a spread of δ_i and other dealers using a bid ask spread of δ_{-i} , is the sum of expected payoffs from each subnetwork of these client node types, C_j , $j = 1, \dots, 2^N - 1$ each of which has a probability or weight in the network of $p(\cdot)$ from an independent distribution.

$$E[\Pi_{D_i}(\Lambda)] = M \sum_{i=1}^{2^N-1} p(i) E[\Pi_{D_i}(C_i), N[C_i]] \quad (4.23)$$

As in section 2.4, we use a simplified network structure with preferential selection, generated by parameters (α, β) where the probability of only 1 link was calculated previously, and the payoff function to the peripheral dealers becomes:

$$E[\Pi_{D_i}(\Lambda)] = M \left(\beta \left[\frac{1}{N} - \frac{\alpha}{N-1} \right] E[\Pi_{D_i}(G(D_i))] + (1 - \beta) E[\Pi_{D_i}(G(D_1, \dots, D_N))] \right) \quad (4.24)$$

and the core dealer D_N 's expected payoff becomes:

$$E[\Pi_{D_N}(\Lambda)] = M \left(\beta \left[\frac{1}{N} + \alpha \right] E[\Pi_{D_N}(G(D_N))] + (1 - \beta) E[\Pi_{D_N}(G(D_1, \dots, D_N))] \right) \quad (4.25)$$

Where $E[\Pi_{D_i}(\Lambda)]$ is the expected payoff for D_i from the entire network, Λ and is the same as the payoff functions in section 4.8, and demonstrates that knowledge of the network generating parameters (α, β) is sufficient to form an expectation of the payoffs. The exact knowledge of the network is not strictly necessary as long as dealers cannot price discriminate.

As before, each of these 2 subgames is a distinct auction with an associated expected payoff function derived in section 2.7.1 □

Dealer i forms an expectation, given the probability mass function of the network, of all the weights in the network, then sums all these payoff expectations of each subgame (distinct client neighbourhood sets) to get an expected payoff.

The objective for each dealer is to maximize their payoff functions, described by equations, 4.24, 4.25. The existence of an equilibrium point is given in the following section 4.10.1 and a numerical methodology to find it is given in section 4.10.4.

4.10 Equilibrium Pricing Solutions

4.10.1 Existence

Proof of existence of a pure strategy equilibrium can be accomplished by showing that an equilibrium exists for a subnetwork of 1 client and arbitrary numbers of dealers and that this then can be extended to an irregular network by standard theorems.

4.10.2 Case 1 - the regular network

The regular network is defined by each client having the same number of links and each dealer is connected a similar subnetwork of clients. This network can be generated from the α, β parameters by setting $\beta = 1$ and $\alpha = 0$ and fixing the client link number k , so that the resulting network is a k -regular bipartite network. The competition process with each client can be viewed as a first price sealed bid auction with unknown common values. This problem was first analyzed by Wilson (1969), extended by Milgrom and Weber (1982). Suppose there are N dealers and M clients. Each client connects to k dealers randomly. The expected links of each dealer comprise of $\frac{Mk}{N}$ distinct client links and the expected payoff to this dealer, D_i is :

$$E[\Pi_{D_i}] = \frac{Mk}{N} E[\Pi_{D_i}(G(C_j, N[C_j]))] \quad (4.26)$$

Each subgame $E[\Pi_{D_i}(G(C_j, N[C_j]))]$ comprises a first price auction with client C_j and $k - 1$ other dealers, where the client can be viewed as the auctioneer and the dealers the bidders.

The solution in terms of bid ask spreads is given by both Wilson (1969) who gave a general partial differential equation of the bidding functions for the duopoly case and Thompson (2005) who provided an explicit solution in the common value linear bidding case. I have restated their common value auction results substituting a semi bid ask spread s instead of bidding functions as:

Theorem 1. *(Wilson (1969) Thompson (2005)) Suppose we have an OTC market as described above with a set of k homogeneous dealers, each receiving a signal $V_i = V + \epsilon_i$ Where ϵ_i is drawn from a distribution with zero mean and SD σ . For any (continuous) distribution of the signal errors with density f_σ and distribution F_σ there exists a unique Nash equilibrium spread strategy δ^* , where each dealer's bidding strategy is of the form $B(V_i) = V_i + \delta_i$. Suppose the client sets no reservation price and trades at the best price that is observed. The Bayesian Nash equilibrium semi spread for each dealer using the auction results of Wilson is:*

$$\delta^*(k, \sigma) = \frac{\frac{1}{k(k-1)} + \int_{-\infty}^{\infty} x F_{\sigma}^{k-2}(x) f_{\sigma}^2(x) dx}{\int_{-\infty}^{\infty} F_{\sigma}^{k-2}(x) f_{\sigma}^2(x) dx} \quad (4.27)$$

The equilibrium not only exists, it has an explicit formulation. However, as reservation prices and network differences are added, the solution no longer has such an analytical form.

4.10.3 Case 2 - the irregular network

The irregular network has an added complication. Each dealer may be involved in multiple auctions, each of which no longer has the same form as in the regular case. For example, a dealer might be connected to 1 client who only connects to them - a monopoly market and also to a client that is quoting prices with many other dealers in the network (competition). In order to justify the existence of the equilibrium in this network we state 2 simple propositions.

Lemma 1. *Suppose there are two separate auctions A_1 and A_2 for a homogenous product with unknown value V in both auctions. Let signals and bidding functions be as before. Let the set of bidders in A_1 be denoted $X = \{a_1, \dots, a_k\}$ with equilibrium point $(\delta^*, \dots, \delta^*)$ and in A_2 by $Y = \{b_1, \dots, b_l\}$ with equilibrium point $(\beta^*, \dots, \beta^*)$, then if there is a bidder x that is in both X and Y , and this bidder seeks to maximize their total expected payoff across both auctions using a single bid ψ in both auctions, $E[\Pi(A_1|\psi) + \Pi(A_2|\psi)]$, then both auctions A_1 and A_2 still have an equilibrium point.*

Proof. Both A_1 and A_2 are single common value first price auctions and therefore have a Nash equilibrium point by the existence theorems of Wilson (1969), and in our case, we assume a linear bidding function, therefore it is both continuous and differentiable. It can be shown that each bidder has a (unique) best response to any actions of the other bidders (the payoff functions have a single maximum), which means they also have a best response if one player does not play (locally) optimally. Fix actions of other players, then x has a best response to these actions (payoffs are concave and sum of concave functions is concave). Therefore A_1 and A_2 both still have a Nash equilibrium point. \square

Proposition 5. *The irregular network OTC market as described above has a Nash Equilibrium point $(\delta_1, \dots, \delta_N)$*

Proof. This is a consequence of the above Lemma. Suppose we have a network market as described above. Each separate sub-network can be viewed as a separate auction which has an equilibrium point. If one of the dealers is involved in another auction, then the remaining dealers adapt their best responses to account for the sub-optimal choice of this dealer (in the sense that the choice is different from the local nash equilibrium) and by lemma 1 also has an equilibrium point. This reasoning then follows iteratively : if another dealer then also plays locally sub-optimally the remaining dealers still have a best response and so this subnetwork also has an equilibrium point. Therefore the entire network has an equilibrium point. \square

These equilibriums can only be solved numerically, except in the regular network case.

Alternatively, the existence of a pure strategy Bayesian equilibrium to the network game exists follows from the following theorem, reported by Asu Ozdaglar (MIT) (2010) and is a result obtained separately by Debreu (1952), Glicksberg (1952), and Fan (1952):

Theorem 2. *(Debreu (1952)) Consider a Bayesian game with continuous strategy spaces and continuous types. If strategy sets and type sets are compact, payoff functions are continuous and concave in own strategies, then a pure strategy Bayesian Nash equilibrium exists.*

It is possible to show that the payoff functions satisfy these conditions and that they are also continuous and concave by showing that any local maximum is also a global maximum. Graphically it appears obvious for any combination of parameters. Fix δ_{-i} of the other dealers to obtain the payoff wrt own spread choices. See figure 4.8.




FIGURE 4.8: Dealer payoff function wrt spread choice, keeping other dealer spreads fixed

4.10.4 Numerical Solution of equilibrium

The equilibrium solution set of bid-ask spreads is simply a matter of jointly maximizing the N payoff functions of the set of dealers, given by the equations:

$$E[\Pi_{D_i}(\Lambda)] = M \sum_{i=1}^{2^N-1} p(i) E[\Pi_{D_i}(C_i), N[C_i]] \quad D_i = D_1, \dots, D_N \quad (4.28)$$

Following the work of [Krawczyk and Uryasev \(2000\)](#) we can compute an equilibrium for the network game numerically. Their method is briefly described here. We have $i = 1, \dots, n$ dealers participating in the OTC game as described above. Each dealer can take an individual action of δ_i in the Euclidean space \mathbb{R}^+ . All players combined can take a collective action, which is the vector $\Delta = (\delta_1, \dots, \delta_n) \in \mathbb{R}^+ \times \dots \times \mathbb{R}^+$. Let $\Pi_i : \delta_i \rightarrow \mathbb{R}$ be the payoff function of dealer i , choosing the action δ_i . Let $\Delta^1 = (\delta_1, \dots, \delta_n)$ and $\Delta^2 = (\zeta_1, \dots, \zeta_n)$ be 2 action sets contained in the action space then $(\zeta_i | \Delta^1)$ is defined as $(\delta_1, \dots, \delta_{i-1}, \zeta_i, \delta_{i+1}, \dots, \delta_n)$ which is the action set where the i^{th} dealer 'tries' ζ_i while the remaining dealers play $\delta_j, j = 1, 2, \dots, i-1, i+1, \dots, n$. Let the collective action set be as above and let $\Pi_i : \delta_i \rightarrow \mathbb{R}$ be the payoff functions of the n dealers. A point $\Delta^* = (\delta_1^*, \delta_2^*, \dots, \delta_n^*)$ is called the Nash equilibrium if, for each i ,

$$\Pi_i(\Delta^*) = \max_{(\delta_i | \Delta^*) \in \Delta} \Pi_i(\delta_i | \Delta^*) \quad (4.29)$$

at δ^* no dealer can unilaterally improve their payoff and is a Nash equilibrium.

Let Π_i be the expected payoff function of dealer i . The Nikaido-Isoda function, [Nikaidô and Isoda \(1955\)](#), Ψ is defined as :

$$\Psi(\Delta^1, \Delta^2) = \sum_{i=1}^n [\Pi_i(\zeta_i | \Delta^1) - \Pi_i(\Delta^1)] \quad (4.30)$$

wrt to the action sets defined above. Each summand can be viewed as the difference in payoff for each dealer when that dealer tries a different spread action ζ_i . An action set Δ^* contained in Δ (the total action space) is a Nash equilibrium if $\max_{\delta_i \in \Delta} \Psi(\Delta^*, \Delta) = 0$

The Optimal Response Function, Z , at the point Δ^1 is defined as: $Z(\Delta^1) = \arg \max_{\zeta_i \in \Delta} \Psi(\Delta^1, \Delta)$ which is a function that returns the set of dealers' actions (best response vector) where all dealers unilaterally try to maximize their own payoffs given Δ^1 . Using these concepts, the relaxation algorithm as suggested by Krawczyk et al. (2000) can be summarised as follows:

Suppose there is an initial game state $\Delta^0 = (\delta_1, \dots, \delta_n)$, where each dealer would like to find their maximum payoff and we aim to find the Nash equilibrium. Given the optimal response function $Z(\Delta)$ is single valued (vector), the relaxation algorithm is given by:

$$\Delta^{t+1} = (1 - \gamma)\Delta^t + \gamma Z(\Delta^t) \quad (4.31)$$

where γ is a fixed constant ($0 \leq \gamma \leq 1$). The result is a weighted average of the improvement point $Z(\Delta^t)$ and the current point Δ^t . The convergence to the Nash equilibrium is guaranteed if certain conditions on the payoff functions are met. It is sufficient for the payoffs to be concave functions in own strategies (i.e have a single maximum and no local maximums). I will just state without proof that the conditions are met in our dealer auction game. Specifically, we numerically find the equilibrium for each dealer by iteratively finding the best response of each dealer to the current dealer action set. The sequence of spreads adopted by a dealer i at step $(t+1)$ is:

$$\delta_i(t+1) = \gamma \delta_i(t) + (1 - \gamma) \arg \max_{\delta_i} E[\Pi_i(\delta_i, \delta_{-i}(t))] \quad (4.32)$$

for each i in turn and for some step size γ (typically 0.7 is a reasonable choice). It is an iterative sequence of best responses to other dealers' best responses or alternatively a dynamic hill climbing optimization algorithm (the hill changes slightly after each

step).

See appendix 4.10.5 for an evaluation of the algorithm convergence to the equilibrium solution and appendix B.0.2 for a comparison with an Agent Based Model solution. The algorithm converges to the equilibrium point in approximately 15-20 iterations of the algorithm, even in quite complex networks.

4.10.5 Algorithm Convergence to NE

As a test for the algorithm, we can see how the solution converges to a known solution. We can use the solution to the common value auction optimal linear bidding strategies as calculated by Wilson (1977)). Suppose we have an OTC market as described above with a set of n homogeneous dealers, each receiving a signal $V_i = V + \epsilon_i$ Where ϵ_i is drawn from a distribution with zero mean and SD σ . For any (continuous) distribution of the signal errors with density f_σ and distribution F_σ there exists a unique Nash equilibrium spread strategy δ^* . Suppose the client sets no reservation price and trades at the best price that is observed. Then expected dealer i payoff is,

$$E[\Pi_i(\delta_i, \delta_{-i}, \sigma_M, \sigma_T)] = - \int_{-\infty}^{\infty} (t - \delta_i) \left[\prod_{\substack{j=1 \\ j \neq i}}^k F_{\sigma_M}[t + \delta_j - \delta_i] \right] f_{\sigma_M}(t) dt$$

and the Bayesian Nash equilibrium semi spread for each dealer is:

$$\delta^*(n, \sigma) = \frac{\frac{1}{n(n-1)} + \int_{-\infty}^{\infty} x F_{\sigma}^{n-2}(x) f_{\sigma}^2(x) dx}{\int_{-\infty}^{\infty} F_{\sigma}^{n-2}(x) f_{\sigma}^2(x) dx}$$

Now consider the network comprising of M clients, each connects to all n dealers. This is a complete bipartite market network and we group client games as before

The payoff from this game for each dealer is equivalent to the payoff in a single client network as above, and so we can check the algorithm against the algebraic expression for the Nash equilibrium. We could also use a symmetry argument for the M clients and the N dealers. Let $n=5$ and assume a standard normal distribution of errors (with mean 0 and sd 1) and indeed, the algorithm finds the theoretical Nash equilibrium quickly, which is:

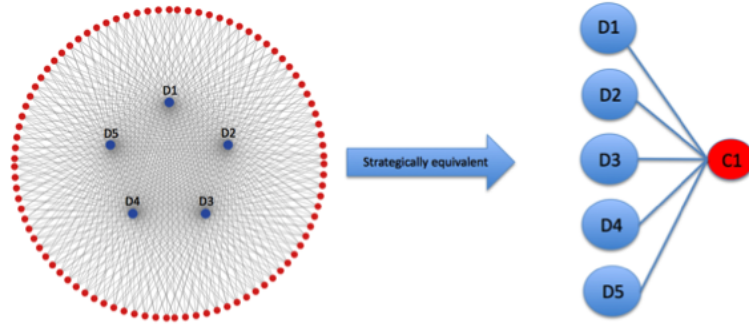
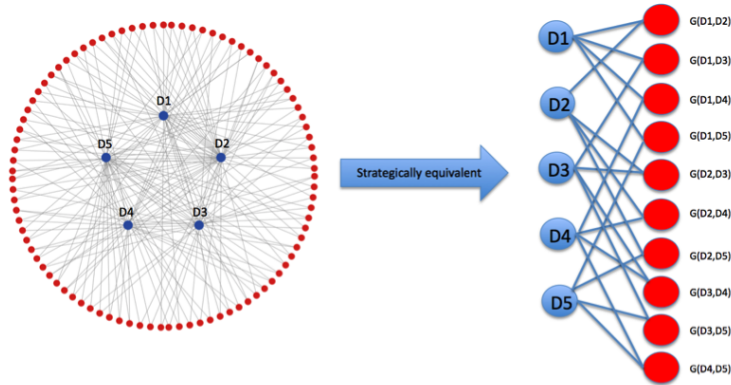


FIGURE 4.9: Strategic equivalence of networks

$$\delta^*(5, 1) = \frac{\frac{1}{20} + \int_{-\infty}^{\infty} x F_1^3(x) f_1^2(x) dx}{\int_{-\infty}^{\infty} F_1^3(x) f_1^2(x) dx} = \frac{\frac{1}{20} + \int_{-\infty}^{\infty} x \Phi_1^3(x) \phi_1^2(x) dx}{\int_{-\infty}^{\infty} \Phi_1^3(x) \phi_1^2(x) dx} \approx 1.5478$$

This is the spread that each dealer adopts in equilibrium. From the clients' perspective, the observed spread Δ , in equilibrium when $n=5$ is: $E[\Delta] \approx 0.7696$

A slightly more complicated market: M clients, all with degree 2 who connect to the 5 dealers randomly. The expected network has M clients with degree 2 and 5 dealers with degree $2M/5$. Grouping the clients with same game structure together gives a regular sparse network:

FIGURE 4.10: Network with client grouping, $G(D_i, D_j)$, represents game G played with dealers D_i and D_j

For each subgame, $G(D_i, D_j)$, the NE formula above can be used to give a NE spread of ≈ 1.7725 (≈ 1.2083 , effective spread to client) and since it is the same NE spread in all subgames, it must be the NE spread for the whole network. Again the algorithm correctly converges to the theoretical NE:




FIGURE 4.11: Algorithm convergence to Nash equilibrium in the 2 networks, $N(0,1)$ signal errors

Chapter 5

Equilibrium Pricing in the Network

5.1 Introduction

Traditional market microstructure models are designed for liquid markets, where dealers can offset trades by adjusting their bid and ask prices. Unfortunately, OTC markets are characterized by low liquidity and relatively large trade sizes and these models cannot totally explain the range of empirically observed pricing anomalies and in particular, the Centrality Premium Effect, which has been documented in many OTC financial markets. [Hollifield et al. \(2012\)](#) find a negative relationship between bid-ask spreads and dealer centrality in the US securitization market. In contrast, [Di Maggio et al. \(2017\)](#) and [Li and Schürhoff \(2019\)](#) find a positive relationship between bid-ask spreads and dealer centrality in the US corporate bond and US municipal bond market.

The centrality premium problem illustrates the limitations of standard market microstructure models and assumptions. In financial OTC markets, the effects of competition, asymmetric information sets, and network topology interact in the price formation process. The microstructure of the market is greatly affected by the topology of who is connected to who. We focus on the observed core-periphery network structure and the link distribution of the clients to dealers and find that these features interact in non-linear ways, producing effects that are magnified by any asymmetric information.

We also examine the effects of increasing transparency using our network model,

where transparency is modeled as a refining of the clients' information sets regarding the asset's true value. We find that the clients' degree distribution is critical in the direction of the effect on spreads. Clients (or markets) with low connectivity (a degree of less than 3) benefit from increased transparency through lower bid-ask spreads. In contrast, clients with a higher degree suffer increased spreads in equilibrium. This is due to dealers raising spreads in equilibrium, to compensate for the increasing risk of the winners' curse which is increasing in the relative information sets of clients and dealers. This result might offer some explanation as to the mixed results when regulators have mandated increased transparency.

Finally, we apply the model to apparently unconnected empirical puzzles, such as why do safe municipal bonds have bigger spreads than much riskier bonds or equities, and find the answers are very related to the same fundamental drivers of the centrality premium, namely, heterogeneous client link formation, preferential selection, and asymmetric information. The effect of illiquid bonds having tighter spreads than very liquid bonds has been documented by Goldstein et al. (2007) from an analysis of the US corporate bond market spreads and liquidity. Our model shows how these markets' topological and information asymmetries are consistent with this effect.

5.2 Bid-Ask Spreads and Uncertainty

Our model predicts that bid-ask spreads will rise with uncertainty over the final asset value. As the asset becomes increasingly hard to measure, bid prices will fall and ask prices will rise.

Our model's equilibrium in a regular network is linearly related to the standard deviation of this measurement error. In fact, if δ^* is the equilibrium dealer bid-ask spread in a market with M clients all connecting to N dealers, and each receiving a signal with standard deviation $\sigma_T = \lambda\sigma_M$ and σ_M , with λ an asymmetry coefficient, then if they receive signals $a\sigma_T$ and $a\sigma_M$, $a > 0$, the equilibrium dealer spread will be $a\delta^*$.

This can be seen in the set of payoff equations, for dealers $i=1,...,N$ at the equilibrium:

$$E [\Pi_{Di} (G(D_1, ..., D_N))] = - \int_{-\infty}^{\infty} \frac{s}{\sigma_M} \phi \left(\frac{s + \delta^*}{\sigma_M} \right) \left[\prod_{j=1, j \neq i}^N \Phi \left(\frac{s + \delta^*}{\sigma_M} \right) \right] \Phi \left(\frac{s}{\lambda \sigma_M} \right) ds \quad (5.1)$$

Set $\sigma_M = a\sigma_M$, and $\sigma_T = a\lambda\sigma_M$, where λ is an asymmetry coefficient, giving a new set of payoff equations:

$$E [\Pi_{Di} (G(D_1, ..., D_N))] = - \int_{-\infty}^{\infty} \frac{s}{a\sigma_M} \phi \left(\frac{s + \delta_i}{a\sigma_M} \right) \left[\prod_{j=1, j \neq i}^N \Phi \left(\frac{s + \delta_j}{a\sigma_M} \right) \right] \Phi \left(\frac{s}{a\lambda\sigma_M} \right) ds \quad (5.2)$$

Substitution $a\delta_i$ for δ_i , then a simple change of the integration variable from s to aS ,

$$E [\Pi_{Di} (G(D_1, ..., D_N))] = -a \int_{-\infty}^{\infty} \frac{s}{\sigma_M} \phi \left(\frac{s + \delta^*/a}{\sigma_M} \right) \left[\prod_{j=1, j \neq i}^N \Phi \left(\frac{s + \delta^*/a}{\sigma_M} \right) \right] \Phi \left(\frac{s}{\lambda \sigma_M} \right) ds \quad (5.3)$$

Leads to the same equation as equation 5.1 (with a higher constant multiplier a , signifying higher profits). This equation 5.1 has an equilibrium point of δ^* , therefore equation 5.3 has an equilibrium of $a\delta^*$. This result shows that the equilibrium bid-ask spreads scale linearly with standard deviation and therefore scale to the square root of time in a multi-maturity asset.

Another implication of this result is the standard observation that reducing asset value uncertainty will reduce transaction costs (and dealer profits) in a market. From a dealers' incentive perspective, they should dedicate resources in the hardest to value complex instruments with the most informational asymmetry. Complex derivative securities probably fall into this category.

An empirical example of this time uncertainty effect can be seen to occur in the US Treasury market for increasing the maturity of bonds. As the bond's maturity increases, the factors that go into the pricing become more uncertain (inflation expectations amongst others), and so spreads will rise. This effect was documented by Brandt and Kavajecz (2004), who analyzed the US treasury market and concluded that "Considering the three pieces of evidence together, we are confident that the yield

changes associated with orderflow imbalances are not attributed to liquidity/inventory risk premiums. The evidence is instead fully consistent with (and further supportive of) our hypothesis of price discovery”. Figure 5.1, taken from Brandt, shows the increases in spreads in on-the-run US treasuries with increasing maturity.




FIGURE 5.1: Dealer spreads in US Treasuries

We can use the model to predict the shape of the bid-ask spreads relative to maturity, and we plot these results on the right of figure 5.1. This plot gives the bid-ask spreads of the yield curve normalized at the bid-ask spread for the 0-6 month security.

The bid-ask spreads in our model scale with the square root of time; that is, a security with double the maturity has a bid-ask spread of $\sqrt{2}$ of the spread in the shorter maturity. This scaling is consistent with a measurement uncertainty whose variance is linear in the maturity and consistent with our model assumptions on the estimation of true value. A similar effect appears to be present in these empirical observations.

The US Treasury market has 22 primary dealers that make bid-ask prices for US Treasury securities. This OTC market’s price-setting process, similar to other bond markets, has some similarity to a sealed-bid first-price auction.

5.3 Bid Ask Spreads and Transparency

There are two driving features in our OTC game: first, the agents’ information sets, characterized by their signal error over the true common value, and the competition effect introduced by multiple dealers and their network topology.

Market transparency refers to the ability of market participants to observe information about the trading process, [Madhavan et al. \(2005\)](#) and can be viewed as a potential to refine the agents' information sets by observing prices and trades of other agents. It has been a major focus of financial market regulators, particularly since the financial crisis of 2008/9. In our model, we can examine the changes to quoted price equilibriums by varying the signal variance to the clients and dealers to mimic changes in transparency. We use the representation that clients receive a normal signal with standard deviation σ_T and dealers σ_M but could equally view $\sigma_T = \lambda\sigma_M$, with λ an asymmetry information parameter. With increased transparency of prices, both sets of agents get an opportunity to refine their estimation of the true value by either a Bayesian updating approach or a heuristic methodology (for example, the increasingly popular learning process presented by [DeGroot \(1974\)](#), which uses an averaging process to refine the agents' signals) and these learning processes lead to a reduction in the variance of the estimates.

We will examine two main network configurations: a regular market network where clients have a fixed number of links, connecting equally to the dealers, and a second network with a preferential selection for one of the dealers.

We focus on a five dealer market, and therefore, each client can connect to the dealers with 1 to 5 links. The network with each client having only one link is the monopoly type network, and the network with each client having five links is the complete or fully connected network, where every client is connected to every possible dealer.

We fix the variance of the homogenous dealers' signal error at 1 and vary the signal error of the clients to mimic changes in relative client information sets. This is then equivalent to varying the asymmetry coefficient λ above. We then compute the dealer equilibrium bid-ask spreads using a numerical algorithm for each set of signal variances at each point, by solving the joint maximization problem using the algorithm described in the previous chapter.

5.3.1 Preferential Selection Network Payoffs

As in section [4.8](#), we introduce an element of preferential selection into the dealer network by allowing one of the dealers (D_N) to have an increased probability α of

being selected from pure randomness. Each dealer with 1 link is equally likely to be selected, with probability $1/n$. Suppose dealer N , D_N is preferred and its' probability of being selected is increased by α . The probability p of D_N being selected is $p = \frac{1}{N} + \alpha$ and the other dealers $[D_1, D_2, \dots, D_{N-1}]$ have a probability of $\left(\frac{1}{N} - \frac{\alpha}{(N-1)}\right)$ of being selected. Suppose also that a proportion β of clients have degree 1 and $(1-\beta)$ have degree N . Our expected biadjacency matrix now looks like:

$$A = \begin{bmatrix} \beta \left[\frac{1}{N} - \frac{\alpha}{N-1} \right] & 0 & 0 & \dots & 0 & 0 & (1-\beta) \\ 0 & \beta \left[\frac{1}{N} - \frac{\alpha}{N-1} \right] & 0 & \dots & 0 & 0 & (1-\beta) \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & (1-\beta) \\ 0 & 0 & 0 & \dots & 0 & \beta \left[\frac{1}{N} + \alpha \right] & (1-\beta) \end{bmatrix}$$

We can now use this variable network representation in the payoff function of a dealer in this network game: The non-preferred dealers $\{D_i\}$ have an expected payoff function in the whole network of:

$$E[\Pi_{D_i}(\Lambda)] = \beta \left[\frac{1}{N} - \frac{\alpha}{N-1} \right] E[\Pi_{D_i}(G(D_i))] + (1-\beta)E[\Pi_{D_i}(G(D_1, \dots, D_N))] \quad (5.4)$$

and the preferred dealer D_N has an expected payoff:

$$E[\Pi_{D_N}(\Lambda)] = \beta \left[\frac{1}{N} + \alpha \right] E[\Pi_{D_N}(G(D_N))] + (1-\beta)E[\Pi_{D_N}(G(D_1, \dots, D_N))] \quad (5.5)$$

Where $E[\Pi_{D_i}(G)]$ is the expected payoff of D_i from the game G .

And the exact payoff in each subgame is given by equation 2.10 in section 2.7.1 giving:

$$E[\Pi_{D_i}(G(D_i))] = - \int_{-\infty}^{\infty} \frac{s}{\sigma_M} \phi\left(\frac{s + \delta_i}{\sigma_M}\right) \Phi\left(\frac{s}{\sigma_T}\right) ds \quad (5.6)$$

$$E[\Pi_{D_i}(G(D_1, \dots, D_N))] = - \int_{-\infty}^{\infty} \frac{s}{\sigma_M} \phi\left(\frac{s + \delta_i}{\sigma_M}\right) \left[\prod_{j=1, j \neq i}^N \Phi\left(\frac{s + \delta_j}{\sigma_M}\right) \right] \Phi\left(\frac{s}{\sigma_T}\right) ds \quad (5.7)$$

Equations 5.4, 5.5, (together with 5.6 and 5.7) give the exact payoffs for all dealers in the OTC network game using a Bertrand price competition with an unknown common value asset. Our aim now is to jointly maximize these payoffs to find an equilibrium solution of the dealer action set of bid ask spreads $(\delta_1, \dots, \delta_N)$.

5.3.2 The Regular Network

We define the k -regular network as a bipartite network where each client has k links and connects to the dealer set equally. Since all of the dealers are connected to the same client subsets, the equilibrium spread is the same for all dealers. To recap, the payoff to dealer i , in competition with $N-1$ other dealers is:

$$E [\Pi_{Di} (G(D_1, \dots, D_N))] = - \int_{-\infty}^{\infty} \frac{s}{\sigma_M} \phi \left(\frac{s + \delta_i}{\sigma_M} \right) \left[\prod_{j=1, j \neq i}^N \Phi \left(\frac{s + \delta_j}{\sigma_M} \right) \right] \Phi \left(\frac{s}{\sigma_T} \right) ds \quad (5.8)$$

Where here, we calculate for each N the number of dealers from 1 to 5, setting $\sigma_M = 1$, varying σ_T and solve for the equilibrium.

The results are illustrated in Figures 5.2 and 5.3 which shows how, in equilibrium, as the relative signal error of the clients rises, dealer spreads reduce up to the point that clients and dealers are roughly equally well informed (similar to the Glosten Milgrom microstructure results). After this point, the dealer spreads progressively increase in the 1 and 2 client degree networks but progressively decrease in the 3, 4, and 5 client degree networks. Two main features drive the equilibrium pricing. Firstly, the information asymmetry causes spreads to rise - for protection against the winner's curse in better-informed clients and for increased profits in the case of worse-informed clients. Secondly, the amount of competition causes spreads to fall. The speed of both of these opposite effects is different for different client degrees, causing the curve shapes in Figures 5.2 and 5.3.

The intuition for this is that with only 1 or 2 dealer connections, the spread lowering effect of competition is very dominated by the spread increasing effect of the increasing relative signal error. In contrast, in the higher degree networks, the beneficial competition effect of the extra dealers quickly dominates the client's lack of information.

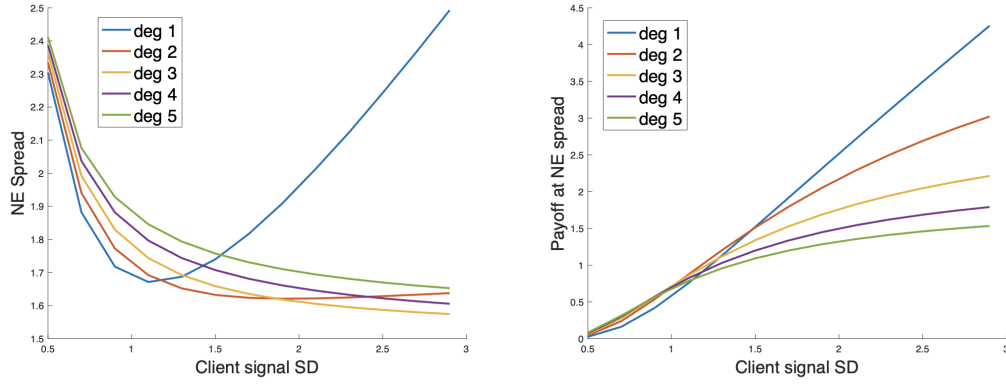


FIGURE 5.2: Dealer spreads to degree and asymmetric information

As client signal error becomes relatively greater, clients make more 'mistakes,' and it becomes increasingly profitable to trade with them, therefore in equilibrium, the dealer spreads tend to reduce. The behavior of 1 and 2-degree networks with respect to asymmetric information is markedly different to 3 or more-degree networks. It appears that a minimum of 3 competitors is needed for the competition effect to dominate asymmetric information; if clients do not set reservation prices, a minimum of 4 dealers is required for the competition effect to dominate any client-dealer asymmetric information.

Bessembinder et al. (2020) analyzed the corporate bond market after the introduction of the TRACE reporting system of increased transparency. He finds a small spread cost advantage to large traders, and a much larger spread cost advantage to smaller customers. Although the network configuration is not discussed, it is reasonable to assume that larger traders' degree is greater than that of smaller traders. Again, the empirical paper by Hendershott et al. (2020) supports this assumption in the UK corporate bond market. In addition, the corporate bond-buying program in the UK, as part of the QE mandate, caused all corporate bonds' prices to rise. However, the bid-ask spreads of those sterling corporate bonds eligible for the program narrowed by 5.3 basis points more than the ineligible bonds after the announcement. Bank of England report on the effects of the QE program on corporate bonds (2018). This QE announcement effectively reduced the uncertainty over these eligible bonds' value, and our model suggests that a reduction of bid-ask spreads would accompany this.

Increasing Transparency		
Client Degree		NE Spreads
1		Decrease
2		Decrease
3		Increase
4		Increase
5		Increase

FIGURE 5.3: Effects of increasing information transparency with client degree when clients have high signal error

The effect of increasing transparency in asymmetric information networks is apparent and summarized in Table 5.3, assuming that clients are initially much worse informed than the dealers. Decreasing client signal variance, which can be viewed as a proxy for increasing transparency, reduces spreads to the lowest degree clients (degree 1 and 2) until clients become as well informed as the dealers. At this critical point, spreads rise again (similar to Glosten and Milgrom (1985) and other informed trader models), but for those clients with a higher degree (> 3), the decreasing signal variance translates into higher spreads from the dealers regardless of their signal variance. The important feature here is that increasing transparency (reducing relative client signal error) is a non-linear effect - it can reduce spreads for the least connected clients (2 or fewer links) but can increase it for the most connected clients. The client degree distribution in the network will determine the effect of increasing transparency on overall market spreads. This transition between degree 2 and degree 3 can be seen more clearly in figure 5.4 by extending the client asymmetry and removing the degree 1 curve.

5.3.3 Bipartite Preferential Attachment Network

We can extend this analysis by examining the effects of preferential selection of certain dealers on bid-ask spreads and the resulting network structure. We have the same 5-dealer network, but now clients connect to each dealer with equal probability and to a core dealer with a probability related to the variable α , which is the excess probability from random chance that a dealer is selected. $P(\text{central dealer is selected}) = \frac{1}{5} + \alpha$, $0 \leq \alpha \leq \frac{(n-1)}{n}$. The dealers now form two groups - the central dealer and the peripheral dealers. Critically, we have non-homogenous client degree - there are 50%




FIGURE 5.4: Dealer spreads to degree and asymmetric information

low connectivity (degree 1) clients and 50% high connectivity (degree 5) clients. This simplification is a model of the common degree distribution found in OTC market networks. For instance, in the UK corporate bond market, the average degree is 3.5, BOE (2019)). Furthermore, one-third have a degree one, [Hendershott et al. \(2020\)](#). In the UK betting market, 44% have degree one with a mean degree of 2.7, Gambling Commission (2019).




FIGURE 5.5: Equilibrium Spreads with Asymmetric Information

The dynamics of the equilibrium show similar effects as seen in Figure 5.5 - equilibrium spreads decrease for the central dealers with reducing client signal variance up to a critical point, at which point spreads start to increase as the similarity of the

information sets causes an increase in the adverse selection problem. The peripheral or non-central dealers show a different effect - increasing transparency causes them to increase spreads for high levels of α since they are connected to more high-degree clients. In this example, we numerically calculate the equilibrium spreads for dealers in a market with $\beta = 0.5$ (50% low degree clients, 50% high degree clients) in varying amounts of central-peripheral effect, α . Dealers receive an $N(0,1)$ signal error, and clients receive a varying signal error. As it becomes more central (popular), the core dealer raises its spread faster with increasing client signal error than the periphery dealers. This U-shaped spread effect is due to the clients' non-uniform link distribution - the dealers are balancing the spreads to clients in low and high competition sub-games. These non-linear effects in transparency changes are critically dependent on the extent of client information asymmetry. The extent of the central-peripheral structure compounds the effect, and these effects lead naturally to the centrality premium puzzle.

These results offer insight into the mixed results when regulators have tried to increase transparency in financial markets. The empirical literature on transparency mainly agrees that increases in transparency reduce bid-ask spreads - it also tends to find that average client connections are also low, for example, the Bank of England report on corporate bonds¹ find client degree around 3 and Municipal Bonds in the US at less than 2. Biais (2005), in his paper summarizing some empirical findings of bid-ask spreads in OTC financial markets, noted that Edwards et al. (2007) had found that increasing transparency in US corporate bond markets reduces spreads. Additionally, increasing transparency in the New York stock exchange in 2001 led to a reduction in bid-ask spreads (see Boehmer et al. (2005)). In contrast, when transparency was increased on the Toronto stock exchange, this caused spreads to increase (see Madhavan et al. (2005)).

Two features determine the effect of increasing transparency - the initial state of the client's relative informedness and the client link distribution. Our model suggests that if clients are less informed than the dealers, increasing transparency will reduce spreads in low degree networks; however, if the clients are equally able to value

¹Bank of England study into UK corporate bond market found average degree of clients is 3 - Staff Working Paper No. 813 Resilience of trading networks: evidence from the sterling corporate bond market (2019) David Mallaburn, Matt Roberts-Sklar, and Laura Silvestri.

the asset as the dealers, increasing transparency may cause bid-ask spreads to rise independently of the network topology. Since it is reasonable to assume that in most markets, professional trading firms have a better ability to measure the true value of an asset than clients, and clients typically have very few links, we would expect that increasing transparency would reduce most market spreads.

5.4 Centrality Premiums in the Network Model

The centrality premium effect can be described as the tendency for more centrally located dealers in a network to make wider prices than those less centrally located, and has been empirically observed fact in many OTC markets by many authors.

In terms of existing literature, numerous other models predict a centrality discount, including Neklyudov (2013), Weller (2013), and Zhong (2014), but these models do not provide a natural explanation for a centrality premium. In Zhong (2014), a centrality premium can arise unconditionally, but there is a centrality discount conditional on trading volume. Similarly, other models predict a centrality premium and use a trading volume reasoning, for example, Uslu (2015), who provides conditions for both a centrality premium and a centrality discount. A centrality premium arises when core dealers intermediate large enough trades in equilibrium. Another example is by Hollifield et al. (2012), who describe an outside options analysis of the securitization market, which tends to exhibit a centrality discount.

In the empirical results, a centrality premium is found by Di Maggio et al. (2017) for corporate bonds and Li and Schürhoff (2019) for municipal bonds, whereas a centrality discount is found by Hollifield et al. (2012) in the securitization market. Our theoretical model, similar to Hollifield, provides conditions on when a centrality discount or a premium can arise. In contrast to the general outside options analysis of Hollifield, our model tackles the root payoff functions of the dealers in a price competition (a common value auction) game and find that the centrality effect is a consequence of the Bayesian Nash equilibrium pricing strategy in a network auction game when there are some preferential selection and asymmetric information.

The centrality premium effect, which is observed in many OTC markets, is typically characterized by having a large unsophisticated client base, whereas the centrality discount appears to be confined to specialized markets with a more sophisticated client case, as also noted by [Hollifield et al. \(2012\)](#).

In our model, if clients are less informed than the dealers and there exists a preferred dealer (core-periphery structure), then the preferred dealer will have a higher centrality and will make wider prices in equilibrium than other dealers. The preferred dealer feature implies a central/less central or core-periphery dealer network structure, which is empirically observed across most OTC markets.

5.4.1 Stylised Model of OTC Network

Similarly, as before, we can model the effect of a single dealer being more or less popular and central than the other dealers by way of the biadjacency weights and using the matrix derived in section 4.8. The idea is to vary the dealer's attractiveness (the core) relative to the others (periphery) and examine the Nash equilibrium dealer strategies. Similarly, for the clients, we can vary our parameter β , to change the proportion of low-degree clients in the market. This will give some insight as to the sensitivity of the dealer equilibriums to the client link distribution. (as mentioned previously, client links are often distributed in an approximately scale-free distribution)

Consider a 5-dealer market and suppose dealers have equal probability of being selected by the clients and we vary dealer 5 probability of being selected by α . Dealer 5 now has probability $(0.2 + \alpha)$, dealers 1-4 have probability of $(0.2 - \alpha/4)$ of being selected. We vary α between 0 and 0.8, so probability of dealer 5 getting selected by the 1-degree clients, varies between 0 and 1.

Suppose there are β clients with degree 1 and $(1-\beta)$ clients with degree 5. The expected network has a biadjacency matrix:

$$A = \begin{bmatrix} \beta \left(0.2 - \frac{\alpha}{4}\right) & 0 & 0 & 0 & 0 & (1 - \beta) \\ 0 & \beta \left(0.2 - \frac{\alpha}{4}\right) & 0 & 0 & 0 & (1 - \beta) \\ 0 & 0 & \beta \left(0.2 - \frac{\alpha}{4}\right) & 0 & 0 & (1 - \beta) \\ 0 & 0 & 0 & \beta \left(0.2 - \frac{\alpha}{4}\right) & 0 & (1 - \beta) \\ 0 & 0 & 0 & 0 & \beta \left(0.2 - \frac{\alpha}{4}\right) & (1 - \beta) \\ 0 & 0 & 0 & 0 & 0 & \beta(0.2 + \alpha) \end{bmatrix}$$

Where as before, α is a parameter determining the attractiveness of a core dealer, and β is a parameter describing the proportion of low link clients in the network.

Case 1 - Dealers Better Informed than Clients

Relatively better-informed dealers would correspond to the majority of real client dealer markets, where the dealers' analytical and informational resources dominate those of the clients. Setting client signal error to be higher than that of the dealers and varying the parameters, alpha, and beta gives a centrality premium of spreads for the core dealer. See figure 5.6




FIGURE 5.6: Centrality premium sensitivity to various network configurations

As dealer 5 becomes more popular (increases in α), its spreads increase and the other dealers' spreads decrease to compensate for their lack of popularity. The centrality premium, described here as the ratio of the core dealers spread divided by the peripheral dealers' spread, is increasing in the amount of the core dealers' centrality,

α and also increasing in the proportion of low-link clients, β (analogous to the scale parameter in a power distribution).

The centrality premium is a direct consequence of the network properties (scale parameter and preferential selection) and asymmetric information. In other words, heterogeneous clients, dealers, not being equally preferred, and clients that are worse informed than dealers cause this effect. Applying extra transparency to this market




FIGURE 5.7: Centrality premium sensitivity to increasing asymmetric information

may cause the centrality premium to decrease, as presumably, the clients are the main beneficiaries of the extra information. Absolute signal noise levels do not affect it - if transparency has an equal effect on the dealers, then the centrality premium will not change.

Markets with a high proportion of unsophisticated low link clients and a set of preferred dealers will exhibit the highest centrality premiums.

The centrality premium is quite sensitive to asymmetric information between the clients and dealers and is typically increasing in information asymmetry. As a corollary to this sensitivity to the relative signal error, in a period such as the financial crisis of 2008-2009, we would expect that the centrality premium in a given market would increase during this period, where presumably, it was much harder for clients to assess the true asset value accurately.

This effect can be seen in the analysis of [Di Maggio et al. \(2017\)](#), where they analyze the corporate bond market spreads over the financial crisis period and find that central dealer spreads appear to rise more than the peripheral dealer spreads in the uncertain crisis environment - i.e., the centrality premium increases.

Case 2 - Dealers Equally (or Worse) Informed than Clients

Conversely, some OTC markets have a very sophisticated client base, for example, the markets for asset securitization, which have almost no retail clients, [Biais et al. \(2005\)](#). We will model this by having both dealers and clients receive the same signal error, i.e., they are both equally able to accurately value the asset. If clients have the same or better signal noise, we predict a centrality discount:

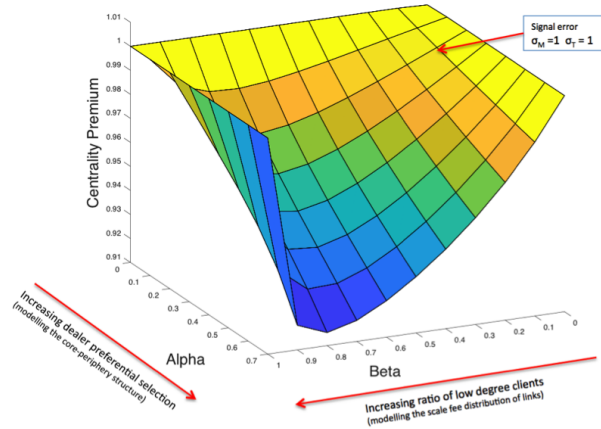


FIGURE 5.8: centrality discount when clients equally informed

Similar to the first case, as dealer 5 becomes relatively more popular, its spreads decrease, and the other dealers' spreads increase. The centrality discount, described here as the ratio of the core dealers' spread divided by the peripheral dealers' spread, is increasing in the amount of the core dealers' centrality, α and also increasing in the proportion of low link clients, β (analogous to the scale parameter in a power distribution).

If clients are equally or better informed, the equilibrium monopoly spread is less than the equilibrium spread charged in competition with multiple other dealers.

An empirical example of this effect was analyzed by [Hollifield et al. \(2012\)](#), and they found that in the securitization market, which is characterized by having almost

no retail clients and only sophisticated investors but still showing a scale-free link distribution, there was a centrality discount in the spreads charged by the dealers.

The absolute centrality effect. (+)ve or (-)ve, is increasing in core dealer centrality and proportions of low degree clients, regardless of the signal error.

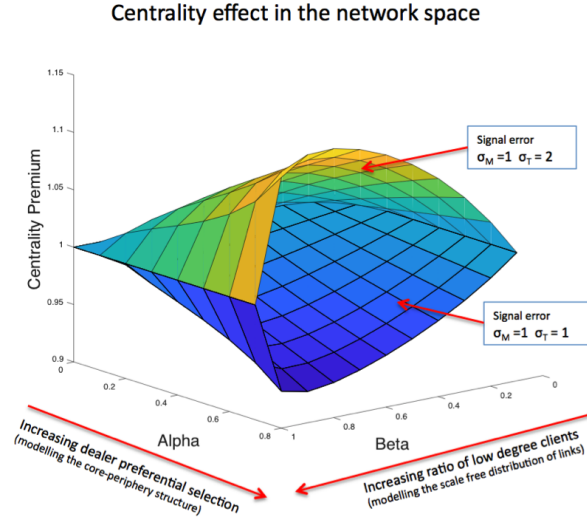


FIGURE 5.9: Envelope of centrality premia with asymmetric information

The envelope of centrality premiums is stretched higher with increases in relative client signal error, and the centrality premium is increasing in the proportion of low-degree clients and the extent of the dealers' preferential selection.

For a given level of core-periphery structure, the centrality premium has a maximum value at approximately 70% low link clients. (The UK domestic energy market, by coincidence, has 70% of users with 1-link and the core suppliers charge much higher prices than the smaller competitors - it appears that the centrality premium for the core suppliers is around 50%. A similar model can be applied to this game, and we would forecast, given the topology, that the core suppliers would have a very large centrality premium in equilibrium).

These premiums and discounts in spreads charged by core and peripheral dealers are common in all markets where there is a degree of preferential selection, and this simple model gives some intuition about the network structures that produce the greatest effects. From a regulatory perspective, the easiest way to combat this centrality premium is to decrease the proportion of low link clients in the network.

5.4.2 1 and 2 Degree Client Case

In a similar manner to the 1 and n-link (monopoly and maximum competition) example outlined above, we can also construct the biadjacency matrix and expected payoff functions for the more messy 1 and 2 link case. Suppose we have n dealers and M clients which have either 1 or 2 links to the dealers. Suppose a proportion $1 - \beta$ clients have degree 1 and β have degree 2. As before each dealer is initially equally likely to be selected with probability $\frac{1}{n}$. Now let dealer n be preferred by an amount α and the probability of being selected by the one degree clients is now $(1 - \beta)(\frac{1}{n} + \alpha)$, with the remaining $n-1$ dealers now having probability $(1 - \beta)(\frac{1}{n} - \frac{\alpha}{n-1})$ of being selected. The probability of the central dealer being selected with one link and another dealer with the other link by a two degree client is $\beta \left((\frac{1}{n} + \alpha) \left(\frac{1}{n-1} \right) + \left(\frac{1}{n} - \frac{\alpha}{n-1} \right) \left(\frac{1}{n} + \alpha \right) \right)$ and the probability of the $(n-1)$ peripheral dealers being paired with another peripheral dealer is $2\beta \left(\frac{1}{n} - \frac{\alpha}{n-1} \right) \left(\frac{1}{n-1} - \frac{\alpha}{n-2} \right)$ where $\beta \in [0, 1], \alpha \in [0, \frac{n-1}{n}]$. These are just the sum probabilities of either being selected by the first link combined with the probability of being selected by the second link, given that you weren't selected from the first link.

In a 5 dealer market with generating parameters (α, β) , the biadjacency matrix becomes:

$$A = \begin{bmatrix} \psi_1 & 0 & 0 & 0 & 0 & \gamma_{12} & \gamma_{13} & \gamma_{14} & \gamma_{15} & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & \psi_2 & 0 & 0 & 0 & \gamma_{12} & 0 & 0 & 0 & \gamma_{23} & \gamma_{24} & \gamma_{25} & 0 & 0 & 0 \\ 0 & 0 & \psi_3 & 0 & 0 & 0 & \gamma_{13} & 0 & 0 & \gamma_{23} & 0 & 0 & \gamma_{34} & \gamma_{35} & 0 \\ 0 & 0 & 0 & \psi_4 & 0 & 0 & 0 & \gamma_{14} & 0 & 0 & \gamma_{24} & 0 & \gamma_{34} & 0 & \gamma_{45} \\ 0 & 0 & 0 & 0 & \psi_5 & 0 & 0 & 0 & \gamma_{15} & 0 & 0 & \gamma_{25} & 0 & \gamma_{35} & \gamma_{45} \end{bmatrix}$$

To recap, the rows represent each dealer. The columns represent the distinct topological connections of the clients and there are ${}^5C_2 + {}^5C_1 = 15$ of these. Columns 1-5 represent the client 1-degree strategies and columns 6-15 represent the 2 link strategies. The entries represent the proportion of the clients that are connected in this way, so $\sum a_{ij} = 1$. For example $a_{11} = \psi_1$ is the proportion of clients with 1 link that select dealer 1, $a_{16} = a_{26} = \gamma_{12}$ is the proportion of clients that select dealers 1 and 2 etc. All of the ψ_i and γ_{ij} are generated by the network parameters (α, β) , that is,

for a give (α, β) , each of the proportions are fixed.

There are general expressions for the values of these biadjacency matrix coefficients and I show the results for the n dealer case. For the $(n-1)$ non-preferred dealers $i=1:n-1$:

$$\begin{aligned}\psi_i &= (1 - \beta) \left(\frac{1}{n} - \frac{\alpha}{n-1} \right) = \Psi_1 \\ \gamma_{ij} &= 2\beta \left(\frac{1}{n} - \frac{\alpha}{n-1} \right) \left(\frac{1}{n-1} - \frac{\alpha}{n-2} \right) = \Gamma_1\end{aligned}\tag{5.9}$$

and the n^{th} preferred dealer:

$$\begin{aligned}\psi_n &= (1 - \beta) \left(\frac{1}{n} + \alpha \right) = \Psi_2 \\ \gamma_{ij} &= \beta \left(\left(\frac{1}{n} + \alpha \right) \left(\frac{1}{n-1} \right) + \left(\frac{1}{n} - \frac{\alpha}{n-1} \right) \left(\frac{1}{n} + \alpha \right) \right) = \Gamma_2\end{aligned}\tag{5.10}$$

The expected payoff function of dealers 1-($n-1$) is the sum of the payoff from the 1 degree clients and the 2 degree clients that they are connected to and is :

$$E[\Pi_{D_i}] = \Psi_1 E[\Pi_{D_i}(G(D_i))] + \Gamma_1 \sum_{\substack{k=1 \\ k \neq i}}^n E[\Pi_{D_i}(G(D_i, D_k))] + \Gamma_2 E[\Pi_{D_i}(G(D_i, D_n))]\tag{5.11}$$

where $E[\Pi_{D_i}(G(D_i, D_j))]$ is the expected payoff to D_i in a price competition with D_j and $E[\Pi_{D_i}(G(D_i))]$ is the payoff to D_i in a monopoly client structure. Obviously $E[\Pi_{D_i}(G(D_j, D_k))] = 0$ for $j, k \neq i$ as this represents the payoff to a dealer in a game they are not involved in.

Similarly, the expected payoff to the preferred dealer, D_n is:

$$E[\Pi_{D_n}] = \Psi_2 E[\Pi_{D_n}(G(D_n))] + \Gamma_2 \left(\sum_{k=1}^{n-1} E[\Pi_{D_n}(G(D_n, D_k))] \right)\tag{5.12}$$

The expected degree centrality of dealer 1:($n-1$) is simply:

$$E[DC(k_i)] = \Psi_1 + (n-2)\Gamma_1 + \Gamma_2\tag{5.13}$$

and the preferred dealer D_n is :

$$E[DC(k_n)] = \Psi_2 + (n-1)\Gamma_2\tag{5.14}$$

for $\beta \in [0, 1]$, $\alpha \in [0, \frac{n-1}{n}]$

These degree centrality measures of the central dealer are strictly increasing in α and β and are plotted in figure 5.10 with $n=5$, a 5 dealer network.




FIGURE 5.10: Raw centrality of central dealer relative to changes in network parameters α, β




FIGURE 5.11: Relationship between α and relative degree centrality $DC(D_n)/DC(D_p)$ of core to peripheral dealers with various proportions of low degree clients

Although both α and β affect the centrality calculation, they affect it in different ways, β affects the global level of network centrality and density, whereas α affects how unequally the links and centrality are distributed. Figure 5.11 shows the centrality of the preferred dealer, normalized by the centrality of the non-preferred dealers.

We can check the equilibriums in this network. All the connections between the clients and dealers are determined by the generating parameters (α, β) . Figure 5.12 graphically summarizes the relative spreads between the central and peripheral dealers in the 1,2 link and 1,n link case. These plots demonstrate that if clients are less

informed (or less sophisticated) to the dealers, the central dealers increase spreads as preferential selection increases. Intuitively, if the clients are equally (or better) informed than the dealers, the central dealer makes tighter prices as preferential selection goes up.

The results are identical to the 1 and n dealer network centrality premium results - as the core dealer relative degree centrality increases, spreads increase.

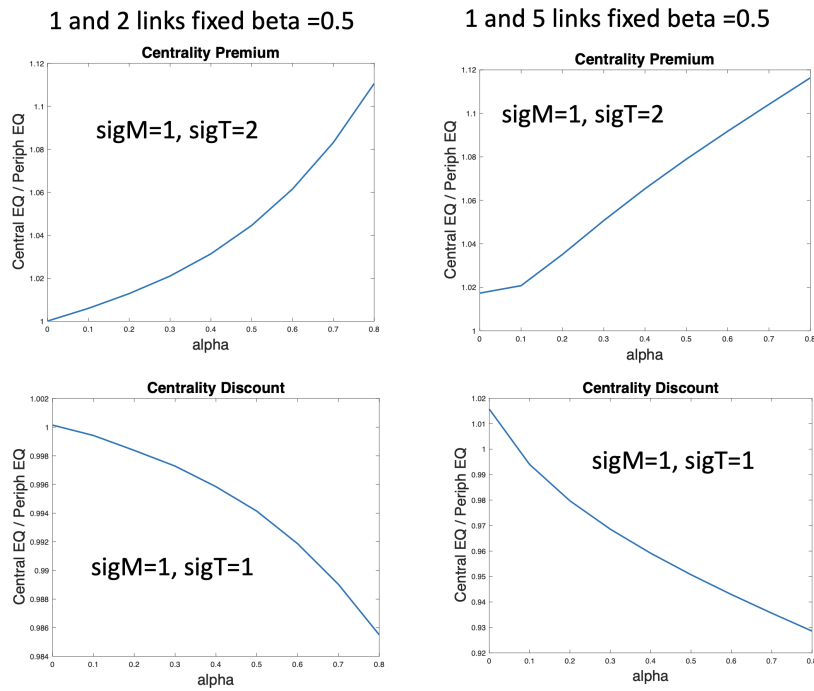


FIGURE 5.12: At the equilibrium, when clients are equally informed there is a centrality discount and when the clients are worse informed there is a centrality premium

5.5 Empirical Anomalies in OTC Markets

5.5.1 Low-Risk and High-Risk Bonds Bid-Ask Spreads

Spreads in low-risk municipal bonds are higher than medium-risk bonds, which are greater still than risky equities in a survey of the empirical literature conducted by [Biais et al. \(2005\)](#). A possible explanation for these counterintuitive observations is likely to be a network effect similar to the drivers of the centrality premium effect discussed above. Municipal bonds have a substantial proportion of retail clients who typically have very few dealer links (the market has a high scale parameter (see [Li](#)

and Schürhoff (2019)). These retail clients would typically have a high signal error and so tend to be charged closer to a monopoly price in a transaction. There are over 2200 broker-dealers in the municipal bond market, and it is infeasible that a retail client would have access to more than a handful. In contrast, institutional clients would be more likely to have multiple dealer-accounts for quoting prices. It is also likely that the retail clients have a higher signal error than the dealers causing an exaggerated centrality premium, as detailed previously combined with high bid-ask spreads. There is also a significant amount of preferential dealer selection, with the top 12 dealers (the core) intermediating over 72% of all transactions.

In summary, high information asymmetry combined with a large proportion of low degree clients causes a large centrality premium and large average bid ask spreads.

According to Biais et al. (2005), "The results obtained by Harris and Piwowar (2004) and Green (2005) are remarkably similar, in spite of the different methods used in the two papers. This convergence of the results speaks in favor of their robustness. Overall, they suggest the Municipal Bond market is highly illiquid, and dealers earn significant markups. Such low liquidity may stem from the lack of transparency of this market. The opacity is such that it is very difficult for retail traders to estimate the market valuation of the security. This puts them in a weak bargaining position."

Conversely, corporate bonds have a much broader blend of client types, hence a higher average client degree, and so the competition element forces the spread down. The core dealers in this market intermediate around 40% of trades, and retail clients account for only 9% of notional volume. This relatively more sophisticated client base causes average degree to rise and relative uninformedness to decline. These two parameters were demonstrated to be primary drivers of both bid-ask spreads and centrality premiums.

In summary, low information asymmetry combined with a large proportion of high degree clients causes a small centrality premium and small average bid ask spreads.

This market structure would tend to suggest high (average) bid-ask spreads and high centrality premiums in the municipal bond market and smaller (average) spreads and smaller centrality premiums in the corporate bond markets. These bid-ask spread predictions were empirically observed by Di Maggio et al. (2017) for corporate bonds and Li and Schürhoff (2019) for municipal bonds.

5.5.2 Liquid and Illiquid Bonds Bid-Ask Spreads

Goldstein et al. (2007) have documented the effect of illiquid bonds having tighter spreads than very liquid bonds in an analysis of the US corporate bond market spreads and liquidity. They found that lower liquidity corporate bonds often had tighter spreads than more liquid issues. Similar to the previous reasoning, the model offers two possible reasons for this: firstly, as above, if low degree clients are more attracted to the liquid issues, spreads will be higher in the liquid bonds, given the same signal error. Alternatively, dealers are very well informed about both sets of bonds' true valuations due to their extensive research resources; however, clients find it harder to measure the illiquid bonds' true value. This leads to a situation where the clients' relative un-informedness is greater in the illiquid bonds than the liquid ones. Our model would predict that in this case, if the clients have a sufficiently high degree, the illiquid bonds would have a tighter spread. Intuitively, this is because clients would typically make more trading 'mistakes' in the illiquid bonds, which leads to higher dealer profits, which leads to lower spreads in equilibrium. Which explanation is more compelling depends on the clients' network properties in these two markets - if clients who trade illiquid bonds quote with multiple dealers, notwithstanding their ill-informedness, then spreads could be relatively low. It illustrates perfectly that the combination of relative information and network properties drives the eventual pricing outcomes.

5.6 Conclusion and Discussion

We have set out a simple theoretical network model of the trading process in a general client dealer market where there is a common unknown asset value, and clients and dealers attempt to estimate the true value of this asset. Clients, either retail or institutional, request quotes from their connected dealers, who then respond with a firm quote to buy and sell, and the client order is executed against the best quote.

Typically, markets are studied either as a monopoly or as a perfectly competitive structure. Our model offers some insight into the interesting behavior that occurs between these two extremes by examining the network topology with the backdrop of asymmetric information sets. By using an auction-based approach to the payoff

functions and finding the dealers' Bayesian Nash equilibrium strategy, we can understand something of the effects of different network configurations and asymmetric information.

As in many market models, for example, the model of [Biais et al. \(2005\)](#), we find the winners' curse problem to be a significant driving force in the determination of dealer spreads and is the main reason why the spreads charged by centrally located dealers can be either a premium or a discount to peripheral dealer spreads. This is an information set phenomenon, but the network effects of client degree distribution and preferential selection play an equally significant role.

We also examine the effects of increasing transparency in these markets, where transparency is modeled as a refining of the clients' information sets regarding the true value of the asset. We find that the degree distribution of the clients is critical in the direction of the effect on bid ask spreads. Markets that have low connectivity clients (a degree of less than 3) benefit from increased transparency by way of lower bid ask spreads, whereas clients with higher degree suffer increased spreads in equilibrium. This is due to the increasing dominance of the winners' curse effect but is also dependent on the initial relative information sets of clients and dealers - relatively uninformed clients benefit most. This could offer a possible explanation for some of the mixed empirical results on the effectiveness of increased transparency on bid-ask spreads.

Finally, we show how our model can offer some explanation to some other puzzles in the empirical literature that are contrary to established economic theory. We saw how the spreads in low-risk municipal bonds can be greater than medium-risk bonds, which are greater still than risky equities. These observations are driven by the client degree distribution and extent of the relative information sets over the true value. Similar reasoning explains the observation that illiquid bonds can have tighter spreads than liquid bonds. These phenomena can also be explained by adding the network topology of preferential selection and degree distribution into the analysis.

Chapter 6

Gambling Markets as Common Value OTC Markets

6.1 Introduction

The UK general election betting market in 2019 offers a unique data opportunity to examine the individual bid-ask spreads in a bilateral market with a common unknown outcome that is likely to be impacted by news events. This market structure is ideally suited to be modeled by an unknown common value price competition over a network. Due to the public nature of bookmakers' individual prices, we were able to collect both the individual bookmaker prices, the best prices in the overall market and exchange prices, and importantly, the betting prices on the range of outcomes from interval level bets (allowing us to create a model of the belief distribution). The betting data was collected over one month leading up to election day.

This chapter's purpose is twofold: first, to understand if the analysis of price formation in the preceding chapters is actually observed in real markets and the extent of any informational content in the betting market prices. Our novel contribution is to look at bookmaker prices with a bid-ask spread price representation in win probabilities and use an unknown common value Bertrand competition methodology in a bilateral network as a model for the trading dynamics between customers and bookmakers.

In terms of betting market microstructure, we can assume a clear partition between bookmakers and customers' functions, allowing a bilateral network representation and their symmetric information sets over the final outcome, modeled by the

equal variance of error signals. There is little chance of any insider knowledge over the final result, and more generally, it is unlikely that anyone has a notably greater knowledge of the final general election outcome. The network of connections between the customers and the bookmakers is sparse (UK Gambling Commission Report); that is, customers only connect to a small number of the total bookmakers and can be represented by a sparse bilateral market network.

I find that, contrary to the standard zero profit assumptions of the full information Bertrand model, bookmakers prices are very close to the predicted Nash equilibrium spreads in an unknown common value FPSBA network model. The regression of predicted values to observed values has an R^2 of 90% with a nearly 1:1 slope between the two with a zero intercept (and well within 95% confidence bounds).

Our model setup of the election gambling market starts with both bookmakers and customers forming a belief as to the true probabilities of each party's chance of winning the election. The bookmakers then post odds prices to their connected customers, and if the implied probabilities of the prices observed are less than the customers' estimation of the true probability, then a unit trade occurs.¹ The bookmakers need to price the initial event probabilities as accurately as possible as it is often difficult that bets can be totally profitably offset.

The odds prices quoted also include some margin or spread so that each bet that the bookmaker lays has a positive expected value. Understanding that his estimate is subject to error is the driving force behind a large part of the bid-ask spread (the bid ask spread is referred to as the overround in betting markets). Indeed, even if the overround were zero, the bookmaker would be expected to lose money over time in competition. This is a function of the celebrated winner's curse – the conditional probability of someone having the wrong price given they were the best price is greater than the unconditional probability of having the wrong price.

We can also compare the bilateral betting market with betting exchanges, which have been a feature of the betting market since 2000. These exchanges are traditional peer-to-peer continuous double auction marketplaces that are open to all participants

¹For example, if the best offered odds price was 2:1 ($a : b = \frac{b}{a+b}$) implying a win probability of $\frac{1}{3}$ and the customer believed the true probability was $\frac{1}{2}$, then they would back the event at 2:1, effectively buying the probability at $\frac{1}{3}$. If the true probability is indeed $\frac{1}{2}$ then they would have an expected payoff of $(\frac{1}{2} - \frac{1}{3} = \frac{1}{6}) \times$ (their stake).

and trade the same or similar products as the individual bookmakers. By far the largest exchange is the Betfair Exchange, which has over 1m customers worldwide and provides bid and ask odds on election betting. Although prices are typically preferable to bookmakers' prices, betting exchanges still account for only a small share ($< 15\%$) of the total online gambling market for reasons that are not well understood. The main differences between the betting exchange and the bilateral market mechanism are; customers can both buy and sell odds on an event, trade with each other rather than just their selected bookmakers, and unwind bets at any time with another counterparty. Most importantly, since opening an account with a bookmaker has a certain cost (time to complete opening procedures and depositing of initial amount), the search costs for the best prices are greatly reduced.

In addition to individual bookmaker data, we also collected the data from this betting exchange, which enables us to compare the two market mechanisms in real-time and answer questions of market efficiency. We notice that although win probabilities are similar, there is a measurable amount of the well-documented longshot bias in the bookmakers' prices compared to the betting exchange – that is, the bookmakers' prices tend to under-price the probability of favorite outcomes and over-price probabilities of longshots. These are similar to the results of [Franck et al. \(2009\)](#), who compared the prices between bookmakers and betting exchanges on football results. They concluded that a strategy that placed bets with bookmakers when their odds were better than the exchange generated positive net returns. This concurs with the idea that the betting exchange provides the most accurate forecasts for unknown events and is consistent with a form of the efficient market hypothesis. We also analyze how the implied probability distribution of voting intentions compares with opinion poll distributions and expert forecasts.

Similar to some OTC financial markets (for example, the securitization market analyzed by [Hollifield et al. \(2012\)](#) also exhibited a centrality discount), we also found a centrality discount in the prices charged by the most popular (central) bookmakers in the gambling market. Although the centrality discount effect is not totally explained using traditional economic theory, our model offers a possible explanation - it is simply a property of a Bertrand competition model when played over a network with preferential attachment.

6.2 Literature Review

The importance of gambling markets for predicting unknown events has been extensively covered by economists and are generally held to be one of the most accurate forms of forecasting in various domains such as sport, politics, and entertainment, [Rothschild \(2009\)](#), [Arrow et al. \(2008\)](#) and others. This forecast accuracy is due to the financial incentives of gathering information and embedding this into the market price, the so-called 'Wisdom of Crowds' effect, popularized by [Surowiecki \(2004\)](#). The Iowa Electronic Markets' election prediction markets have outperformed the accuracy of the large-scale polling organizations, [Berg et al. \(2008\)](#) and Senatorial elections in the US in 2008.

In particular, the literature of using gambling market implied-probabilities in UK elections range from, for example, [Wall et al. \(2012\)](#) which looked at constituency level betting market data to forecast seat shares in the 2010 UK general election, to [Rosenbaum \(1999\)](#) on the accuracy of betting market forecasts in the 1997 UK general election. This paper approaches a less studied aspect of these prediction markets, which is the bookmakers' individual pricing behavior and, in particular, their bid-offer spread and the confidence or belief in their estimates of final value, which can be implied from the interval level bet prices. These interval level bets, which are odds prices for a specific vote share interval, for example, the odds price of a party gaining between 10% and 20%, provide an approximation to the probability distribution of the point estimates of expected final outcomes.

Market microstructure literature would suggest (see [Biais](#) for a comprehensive summary of the microstructure literature [Biais et al. \(2005\)](#)) that bookmakers' odds are expected to move on exogenous news (information effects) about the likely outcome of the event but also on account of the flows of bets that they receive (inventory effects). It appears that inventory effects were not a driving factor of bid ask spreads in the gambling market over the time period (1 month) that we monitored prices and prices moved very little intraday. The conventional bookmakers' strategy has been described by [Wall et al. \(2012\)](#) as 'uses expert knowledge to derive the probability of a given outcome, and then offers customers' odds' at which they can back that outcome.' According to [Levitt \(2004\)](#), "The market for sports gambling is structured

very differently from the typical financial market. In sports betting, bookmakers announce a price, after which adjustments are small and infrequent. Bookmakers do not play the traditional role of market makers matching buyers and sellers but, rather, take large positions with respect to the outcome of game.” This structure is reminiscent of many OTC financial markets such as corporate bonds and derivatives, where liquidity in each individual issue is small, but trade sizes can be quite substantial, ensuring that dealers are incentivized to attempt to price each security as accurately as possible.

6.3 Background

6.3.1 UK General Election Basics

In the UK, General Elections take place every 5 years in May, unless Parliament votes to hold an election sooner, which is what happened to initiate this election in December 2019. Candidates compete for a seat in the House of Commons and the election is comprised of 650 individual elections for each seat or constituency, which each have a similar number of voters. The party that wins a majority of seats, gets to form a government for the next 5 years.

At the start of the election process, there were several important features. Brexit and the consequent importance of the Brexit party vote share were paramount, as it was feared that the Brexit Party votes would hurt the Conservatives disproportionately in its marginal seats. The competitor, Labour, (the UK general election has been won by either Conservative or Labour at every election since 1918 ²) was polling very poorly due to its new leadership and political agenda. This led to initial implied probabilities of 85% for the Conservatives winning the most seats but also a 50% chance of a hung parliament, where no one party has more seats than the sum of the others.

Our analysis focuses on two types of data, the prices (probabilities) associated with the final winner (most seats) and the final vote percentage for each party. The vote to seat problem, that is, the inference of seat share from vote shares, is briefly discussed below.

²www.commonslibrary.parliament.gov.uk

6.3.2 Vote to Seat Prediction Problem

This first past the post system, FPTP, or technically the Single Member Plurality (SMP) electoral system, in 650 individual constituencies complicates the forecasting of the final most seats result from small national samples as each party has core constituencies that require large swings to change the result and other constituencies that need only small swings in the percentage vote to change the winner. Consequently, predicting vote share is a simpler problem compared to predicting seat shares due to the costs of polling representative samples.

The problem of vote share to seat share has been analyzed extensively, and a comprehensive summary of the vote to seat problem is provided by [Whiteley \(2005\)](#). There are three main approaches to the problem. Firstly, the 'cube rule' and various later extensions, which are heuristic rules that are derived from the original empirical observation that the ratio of seat shares to the ratio of the cube of vote shares was approximately equal. This simple rule has been found reasonable success in predicting seat shares where there are two main competitors. Other power-law formulations have been tried with varying levels of success to take account of election-specific idiosyncrasies such as vote fragmentation and geographical clustering. The second set of methodologies is econometric and includes detailed data about previous elections' vote and seat share data. [Lebo and Norpoth \(2007\)](#); [Whiteley \(2005\)](#) both employed time series models that include both a component of previous vote shares, seat shares, and current opinion poll data in an attempt to forecast the predicted seat share. Coefficients were calibrated using all previous UK election data since 1945.

Given the complexities and idiosyncrasies of each election, a perfect general model of seat shares is difficult to achieve; however, the betting markets ask a slightly different but simpler question – what is the probability that a particular party wins the greatest number of seats on the election date. The question reduces the necessity of forecasting seat numbers to which party wins the most seats, which is equivalent to which party has the greatest vote share over a threshold dependent on each vote-share. Figure 6.1 illustrates the nonlinearity of this threshold at a point in time. It shows that, because of the constituency nature of the election, the percentage vote share needed to win depends on both the relative vote-share and the distribution of

the shares.




FIGURE 6.1: Vote Share Thresholds to Win Election
(electoralcalculus.com)

With implied probability density functions of each party's vote share, we can compute winning probabilities of the greatest number of seats without forecasting the seat numbers directly.

6.4 Probabilities and Bid-Ask Spreads from Betting Data

6.4.1 Dataset

Data were collected from 19 bookmakers (16 consistently active) over the period 12 November to 12 December 2019 from published betting prices on the internet – oddschecker.com, a betting aggregation portal, individual bookmaker sites, and the Betfair betting exchange data. Two data types were collected – pricing data of winning the election (defined as winning the most seats) and data indicating the size of the vote share for the three main parties. The odds data were then translated into probabilities to compare bookmakers' actual spread quotes with our theoretical spread predictions. We also compared the implied belief distribution versus the opinion poll

data at the time. Data collection can be difficult as bookmakers deter automated data scrapping from their websites by changing the location of betting prices.



FIGURE 6.2: Sample daily odds prices from main bookmakers

Figure 6.2 shows a screenshot of a typical daily data source from oddschecker.com, summarizing the quoted betting odds from the major bookmakers. Figure 6.6 shows a sample screenshot from a leading bookmaker showing the interval level bets on the percentage vote bands for one of the political parties. These interval level bets allow us to calculate a distribution of beliefs around the mean forecast.

6.4.2 Implied Probability Bid-Ask Calculation

Odds can be easily transposed to probabilities by the following formula. If a bookmaker makes an odds price for the Conservatives (C) to win the election of C1 to C2, then the associated probability is: $P(C \text{ wins}) = \frac{C2}{C1+C2}$. For example, if the quote was 2:1, then the probability of winning is 0.33.

We ignore the very small and nationalist parties and focus only on Conservative, Labour, and LibDem parties (and the Brexit party pre their 11th November announcement of not standing in conservative held seats, after which their expected vote share collapsed to nominal levels).

Bookmakers do not generally provide odds that allow a gambler to lay (sell) a particular bet, but we can implicitly obtain the odds of the bookmakers' 'bid price' odds by simply computing $1 - P(\text{Labour Wins}) - P(\text{LD Wins})$, that is what is,

the probability of conservatives not winning by betting on the other two possible outcomes.

Event	Win Odds	Win Probability Offer	Win Probability Bid
Conservatives	1:20	95.24%	89.92% (implied)
Labour	10:1	9.09%	3.77% (implied)
LibDem	100:1	0.99%	-4.33% (implied)
Totals		105.32%	89.36%

TABLE 6.1: Win and Lose Probabilities derived for Odds Prices

This methodology allows us to view the bookmaker prices as a bid-ask spread pair in probabilities for each event. The quoted win probabilities over all possible events do not sum to 1, and the excess over one is known as the overround. The overround of 5.32% is the excess bookmakers' prot if they laid (sold) each of these bets in equal quantities and is also the expected loss to customers since it is a zero-sum game. Bookmakers can prot from these prices, and in particular the overround, in two ways, firstly if the volume of bets placed is in the exact ratio of the odds is sold (laid), then the bookmaker has a guaranteed prot of 5.32%. i.e., if bookmaker laid £95.24 at 1:20, £9.09 at 10:1 and £0.99 at 100:1, then whichever of the three outcomes occurs, the bookmaker makes a risk-free prot of £5.32. Secondly, for any open bets that the bookmaker has sold, the expected value to the bookmaker of these bets is the difference between the implied probability sold and the true probability multiplied by the volume staked on each bet.

If we assume risk-neutrality and rationality, that is, bookmakers only care to maximize their expected payoff over time, then this implies that each event should have a positive expected value and hence overestimate the true or perceived probability, since the bookmaker cannot be sure what volume will be received for each event. In extremis, a risk-averse bookmaker would adjust prices to maximize the chances of laying the bets in the correct ratio, and the risk-neutral bookmaker would prefer to sell the bet with the highest expected value.

Monitoring the individual bookmaker prices over the month suggests that prices are made with most regard to expected final value as price movements are infrequent and tend to move slowly with the changing expected mean of the distribution of vote shares. See Wang and Pleimling (2019) for a discussion on Wager distributions at the

aggregate level.

The average daily volume of bets split was approximately 48% Conservative, 45% Labour, and 7% others. (source: oddschecker.com, daily % of bets placed through the site). This suggests that the bookmakers did not adjust prices to encourage the risk-free proportion of bets and also that they entered the election with a negative inventory exposure to a Labour win and hence would have been profitable since the conservatives won the election. Prices appeared to move very little on inventory effects and moved purely on new information regarding the outcome, consistent with the bookmaker description given by [Wall et al. \(2012\)](#).

6.4.3 Summary Win Probabilities

The table in figure 6.3 shows the summary results of applying this methodology to the data set and yields the following bid ask spread results for the expected win, defined as the party winning the greatest number of seats, using the individual bookmaker prices and the exchange market as a comparison. The bid ask prices are the highest bid and lowest ask prices across all bookmakers. It should be noted that these are derived from the actual trading prices with no data adjustment for any potential bias:

	0	1	2	2	5	7	12	14	17	21	26	27	28	30	32	33	34	35
	08-Nov	09-Nov	10-Nov	12-Nov	14-Nov	19-Nov	21-Nov	24-Nov	28-Nov	02-Dec	03-Dec	04-Dec	06-Dec	08-Dec	09-Dec	10-Dec	11-Dec	
Conservative	Bilateral	Bilateral	Bilateral	Bilateral	Bilateral	Bilateral	Bilateral	Bilateral	Bilateral	Bilateral	Bilateral	Bilateral	Bilateral	Bilateral	Bilateral	Bilateral	Bilateral	
Best Pr Offer	86.36%	88.89%	89.47%	94.12%	94.12%	95.24%	95.24%	94.12%	95.24%	95.24%	95.24%	95.24%	95.24%	95.24%	95.24%	95.24%	95.24%	
Best Pr Bid	83.21%	85.44%	86.55%	91.32%	91.51%	91.98%	91.98%	91.32%	92.84%	93.00%	93.08%	93.79%	92.93%	93.08%	93.13%	93.13%	93.13%	
Best Pr Spread	3.15%	3.45%	2.93%	2.80%	2.60%	3.26%	3.26%	2.80%	2.40%	2.24%	2.15%	1.45%	2.30%	2.15%	2.10%	2.10%	2.10%	
Labour																		
Pr Offer	13.33%	11.11%	10.00%	7.69%	7.69%	7.69%	7.69%	6.67%	6.67%	6.67%	6.67%	5.88%	6.67%	6.67%	6.67%	6.67%	6.67%	
Pr Bid	10.18%	7.66%	7.07%	4.89%	5.09%	4.43%	4.43%	4.89%	4.26%	4.43%	4.51%	4.43%	4.36%	4.51%	4.56%	4.56%	4.56%	
Pr Spread	3.15%	3.45%	2.93%	2.80%	2.60%	3.26%	3.26%	2.80%	2.40%	2.24%	2.15%	1.45%	2.30%	2.15%	2.10%	2.10%	2.10%	
LibDem																		
Pr Offer	1.96%	1.96%	1.96%	0.99%	0.79%	0.33%	0.79%	0.99%	0.50%	0.33%	0.25%	0.33%	0.40%	0.25%	0.20%	0.20%	0.20%	
Pr Bid	-1.19%	-1.49%	-0.97%	-1.81%	-1.81%	-2.93%	-1.81%	-1.81%	-1.90%	-1.90%	-1.90%	-1.12%	-1.90%	-1.90%	-1.90%	-1.90%	-1.90%	
Pr Spread	3.15%	3.45%	2.93%	2.80%	2.60%	3.26%	2.60%	2.80%	2.40%	2.24%	2.15%	1.45%	2.30%	2.15%	2.10%	2.10%	2.10%	
	08-Nov	09-Nov	09-Nov	12-Nov	14-Nov	19-Nov	21-Nov	24-Nov	28-Nov	02-Dec	03-Dec	04-Dec	06-Dec	08-Dec	09-Dec	10-Dec	11-Dec	
Conservative	Exchange	Exchange	Exchange	Exchange	Exchange	Exchange	Exchange	Exchange	Exchange	Exchange	Exchange	Exchange	Exchange	Exchange	Exchange	Exchange	Exchange	
Pr Offer	87.50%	88.89%	90.00%	94.12%	93.33%	94.12%	94.12%	93.33%	95.24%	95.24%	95.24%	96.15%	96.15%	96.15%	96.15%	96.15%	96.15%	
Pr Bid	86.50%	88.33%	89.29%	92.80%	92.85%	93.89%	93.89%	93.40%	92.69%	94.67%	94.70%	95.43%	95.18%	95.18%	95.14%	95.35%	95.35%	
Pr Spread	1.00%	0.56%	0.71%	1.32%	0.48%	0.22%	0.72%	0.72%	0.64%	0.57%	0.54%	0.73%	0.97%	0.97%	1.02%	0.80%	0.80%	
Labour																		
Pr Offer	11.36%	9.26%	8.47%	6.41%	6.58%	5.81%	5.81%	5.81%	6.58%	5.10%	5.10%	4.35%	4.55%	4.55%	4.76%	4.55%	4.55%	
Pr Bid	10.36%	8.70%	7.76%	5.09%	6.10%	5.59%	5.59%	5.09%	5.94%	4.53%	4.56%	3.62%	3.57%	3.57%	3.74%	3.74%	3.74%	
Pr Spread	1.00%	0.56%	0.71%	1.32%	0.48%	0.22%	0.22%	0.72%	0.64%	0.57%	0.54%	0.73%	0.97%	0.97%	1.02%	0.80%	0.80%	
LibDem																		
Pr Offer	1.35%	1.56%	1.96%	0.79%	0.57%	0.29%	0.29%	0.79%	0.73%	0.23%	0.20%	0.23%	0.28%	0.28%	0.10%	0.10%	0.10%	
Pr Bid	0.35%	1.00%	0.74%	-0.53%	0.09%	0.07%	0.07%	0.09%	-0.34%	-0.34%	-0.50%	-0.70%	-0.70%	-0.92%	-0.70%	-0.70%	-0.70%	
Pr Spread	1.00%	0.56%	0.71%	1.32%	0.48%	0.22%	0.22%	0.72%	0.64%	0.57%	0.54%	0.73%	0.97%	0.97%	1.02%	0.80%	0.80%	

FIGURE 6.3: Summary of Implied Bid Offer Probabilities

The difference caused by the longshot bias is quite apparent in figure 6.3 between the best of the individual bookmaker mid prices compared to the exchange mid prices:

The outsider, Labour, seen in Figure 6.4, is consistently assigned a higher probability (lower odds) by the bookmakers than at the exchange. Similarly, the favourite, Conservative, also seen in Figure 6.4, has consistently lower probability (higher odds)

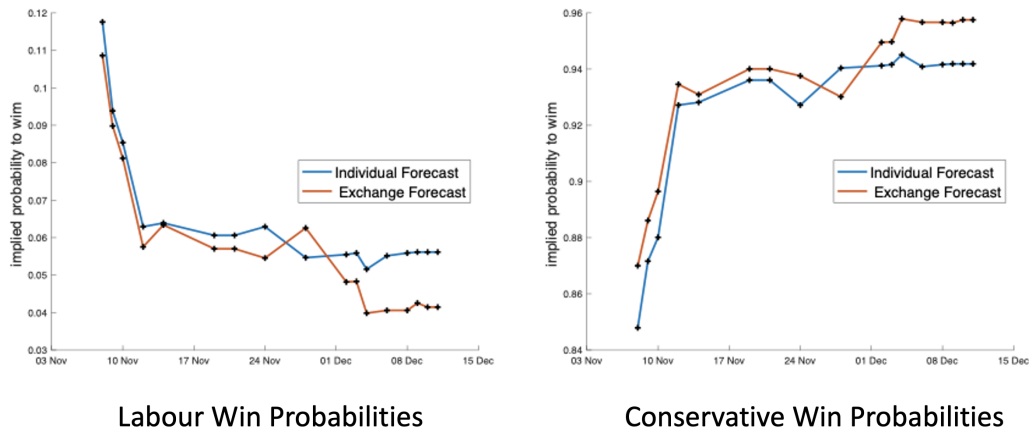


FIGURE 6.4: Average Implied Win Probabilities from Odds Prices

than the exchange using the mid prices. The structural benefit of the exchange allows gamblers the opportunity of adding extra supply to the outsiders (they can lay outsiders directly) and so will tend to mitigate the longshot bias of the bookmakers' prices. The mid-prices suggest that the bookmakers make better (lower) prices than the betting exchange, but this in fact is not true. The larger bid offer spreads with the bookmakers means that the exchange is still a cheaper source of favourite (Conservative) probabilities. This dataset thus allows us to quantify the extent of this bias which I don't believe has been empirically estimated before in election markets.

6.4.4 Bilateral and Exchange Pricing

Unlike most financial markets, the gambling market has both types of market mechanisms trading a homogeneous product – the final election outcome probabilities. In the UK gambling market there are both individual bookmakers making OTC type prices and a betting exchange (the Betfair exchange) with a continuous double auction limit order book open to everyone.

Figure 6.5 shows that spread prices on the exchange are consistently tighter than not only individual bookmakers prices but also the best bid and best offer over all our 16 active bookmakers. This is a well-known phenomenon and was one of the reasons that initially, betting exchanges were thought to be the future of online gambling. For reasons not well understood, this has not occurred and betting exchanges account for < 15% of the online volume.




FIGURE 6.5: Observed Exchange, Bookmaker and Best of Bookmaker Spreads

6.4.5 Adjusting Implied Probabilities from Bias

Inferring the true event probabilities from the probabilities implied by the betting prices suffer from two well-documented problems; the behavioral bias that is known as the longshot bias, which is the tendency for gamblers to under-bet the favorites and over-bet the outsiders and the overround problem, [Peel et al. \(2003\)](#). Although for most of our analysis, specifically, predicted and actual spreads, we use only the quoted prices without adjustment, the calculations for the variance of beliefs require a different method as there are typically nine possible event bands that make the additive method of reducing the overround inaccurate.

The most researched stylized fact of gambling markets is the favorite-longshot bias and was first noted by Griffith in 1949 and has been observed in most non-exchange betting data. Equilibrium market prices are a biased estimate of the true probability of an event occurring, and the expected return to gamblers on longshots is lower than the expected return on favorites at the bookmaker equilibrium prices. See, for example, [Snowberg and Wolfers \(2010\)](#) and [Peel et al. \(2003\)](#) for a complete discussion of this phenomenon.

There are two sets of competing theories to explain the empirical observation of the apparent irrationality of over-betting on the worse-expected value events relative to the best-expected return events. Firstly, a neoclassical approach that explains

this by proposing risk-loving utility functions – the return to gamblers on favorites is small relative to the return from outsiders. Secondly, behavioral theorists suggest an approach that concerns cognitive errors and probabilistic misconceptions. Cognitive psychology has demonstrated that people are systematically poor at differentiating between small and tiny probabilities and treat both similarly [Kahneman and Tversky \(1979\)](#). Odds of 100:1 or 200:1 ($P = 1\%$ and 0.5%) do not affect the demand for that event bet; hence bookmakers prefer to offer at the lower odds, thereby distorting the true win probabilities. Also, people exhibit an irrational preference for certainty over highly probable outcomes, which leads to the under-pricing of favorites - odds of 1:20 and 1:30 (95% and 96% probability) again cause limited demand changes and hence bookmakers tend to prefer to make the 1:30 price. [Snowberg and Wolfers \(2010\)](#) tested a dataset consisting of over 5 million horse races over 1992-2001 and concluded that misperceptions in probability drove the longshot bias.

In the case of the distribution of vote-share beliefs, the longshot bias leads to bookmakers reducing odds for low probability event vote-shares more than for higher probability vote-shares, and the effect of this is that the raw implied probabilities from the odds data will imply a greater variance than the true distribution of beliefs over final outcomes.

The total π of the implied probabilities for each event is known as the booksum and the excess, $\pi - 1$, the overround. Methods have been devised to extract true event probabilities from this overround by normalizing each event's implied probabilities to closer to the true event probabilities. The question of how to translate these betting probabilities into accurate event probabilities has been widely studied, as knowledge of true event probabilities can be used to predict future events, [Vovk and Zhdanov \(2007\)](#) and others.

There are four main methods of translating the actual event probabilities from the betting probabilities and removing the bookmakers' overround (see [Clarke \(2017\)](#) for a full description). Firstly, the Additive Method distributes the overround evenly between the n outcomes as a probability adjustment. The second, and most commonly used method in the literature, is the Normalization Method, which normalizes each implied probability by the booksum, π . Suppose event A has implied probability π_A , this is adjusted to the true probability $P(A)$, by $P(A) = \frac{\pi_A}{\pi}$. This method allocates

the same proportion of the overround to each event and thus does not address the longshot bias. A third method was proposed by Shin, H. S. (1993) and used an iterative method based on an assumed fraction z of informed traders to compensate for the informed traders' effects on prices.

The final method, and the one we use in this paper to normalize the belief probabilities, is known as the Power Method and is originally attributed to Victor Khutishvili and is described by Vovk and Zhdanov (2007) and Clarke (2017). It is the basis for many commercial automated betting algorithms at several bookmakers. The method is a natural extension of the Additive and Normalization Methods and raises each implied probability to a fixed power k in order to reduce the booksum to 100%. $P(A) = \pi_A^k$ with $\sum_A \pi_A^k = 1$. The effect is that the overround is eliminated and affects the low probabilities more than the higher probabilities. This method's predictive power has been empirically demonstrated to be superior to the other three methods after applying it to over 20,000 actual sporting events. This method also appears to compensate for the inherent longshot bias.

6.5 Implied Beliefs and Opinion Polls

6.5.1 Implied Belief Calculation

Some bookmakers provide odds prices for certain percentage vote share bands (known as interval level bets), and these odds can be translated into probabilities and hence an implied distribution with the mean corresponding to the mean belief of the final vote share and distribution around it. See figure 6.6 for a sample data of interval level bets from 888sport bookmaker.

Total Seats obtained by the Party	
Over 38.5	3/4
Under 38.5	19/20
Percentage of votes obtained	
	Winner
4.99% or less	33/1
5.00-9.99%	10/1
10.00-14.99%	11/4
15.00-19.99%	5/4
20.00-24.99%	7/4
25.00-29.99%	7/2
30.00% or more	8/1

FIGURE 6.6: Sample interval level bets from 888Sport bookmaker

This set of contracts reveals an approximation to the full probability distribution of market expectations. Beliefs extracted from the betting market were transformed using the Power method as described by [Clarke \(2017\)](#) because to calculate the implied buying price using the selling prices of all of the other interval level bets requires too many crosses of the bid-ask spread (8 crosses of the spread) and so calculations are not helpful.

Event	Odds	Quoted P(Win)	N-Method P(Win)	P-Method P(Win)
under 14.99%	100:1	0.99%	0.79%	0.47%
15-19.99%	80:1	1.23%	0.98%	0.61%
20-24.99%	40:1	2.44%	1.94%	1.34%
25-29.99%	18:1	5.26%	4.19%	3.28%
30-34.99%	7:1	22.22%	17.69%	17.44%
35-39.99%	7:1	36.36%	28.94%	30.90%
40-44.99%	13:1	38.10%	30.32%	32.61%
45-49.99%	6:1	14.29%	11.37%	10.44%
over 50%	20%	4.76%	3.79%	2.92%
Total		125.66%	100%	100%

TABLE 6.2: Summary of Conservative Party Event Odds and Implied True Win Probabilities using N and P method

Table 6.2 shows a sample data field showing the event and quoted odds by a leading bookmaker on 14th November 2019, together with the normalized probabilities (normalizing each probability by the extent of the overround) and the power normalization that eliminates the overround by raising each probability to a power and somewhat compensates for the longshot bias. The mean and standard deviation are easily calculated from this frequency table as 38.97% and 6.25%, respectively. The mean and variance are then used to fit a normal distribution to the data.

Figures 6.7 represent the histograms of beliefs taken from bookmaker prices on interval level bets of vote share for each party. The probabilities are normalized in the usual way and give an insight into the both the mean and the variance of the point estimates of vote shares.

Graphically, the beliefs data seem well modelled by a normal distribution that is also drawn over the histogram for illustration (using the sample mean and variance). This fitting process is repeated for each bookmaker and also for every day in order to give a daily average belief distribution of the expected vote shares of each party.




FIGURE 6.7: Histograms of Implied Observed Belief Distributions

Goodness of fit of a normal distribution to the data can be confirmed with a standard chi-squared goodness of fit test. Table 6.3 shows the average daily chi-squared test statistics for each party.



	Conservatives	Labour	Libdem
χ^2 statistic	3.97	3.85	3.62
95% Critical Value	9.35	11.14	7.38

TABLE 6.3: Average Test Statistics for Normality of Data

The χ^2 statistics are all comfortably below the critical value and so we are confident that the normal distribution is a good approximation for the belief data which confirms the graphical intuition.

Figure 6.8 plots the mean values of these implied distributions, which not surprisingly, due to non-arbitrage consistency of the odds prices, are consistent with the both the win probabilities and opinion poll data. Figure 6.9 shows the diminishing standard deviation of the belief distributions over time, which is consistent with our modelling assumptions.

As shocks that might substantially alter the perceived outcome fail to occur, mean beliefs trend towards the final outcome monotonically. In addition, the standard deviation (or lack of confidence) in these beliefs decreases in a linear manner as the chances of shocks decrease linearly with time.

FIGURE 6.8: Implied
Vote SharesFIGURE 6.9: Implied
SD of Votes Shares

6.5.2 Opinion Polling and Expected Beliefs

Opinion polls are a perennial feature of most political elections, and 63 opinion polls by 12 opinion pollsters were conducted over the period 1st November – 11th December 2019, with an average of 2 polls published per day. The problem with opinion polls is that they do not exactly measure what we need to know to price a bet on the outcome. The opinion polls aim to measure (with some sample error) the voting intentions on a given day. They do not measure what probability these voters might switch their vote on the actual day of the election due to a variety of factors or the confidence level of each respondent.

The Pew Research Centre has examined how and in which order questions are asked and how these affect the answers given. Determining voter preference among the candidates running for office would appear to be a relatively easy task by simply asking them who they will vote for on Election Day. Differences in how this question is asked and placed in the questionnaire can affect the results. They find that many people have given little thought to the campaign or are genuinely ambivalent about their choices. For these voters, the structure of the questionnaire impacts their answers.

Poll respondents are typically asked a question along the lines of ‘if the general election were tomorrow, how would you vote?’. It is difficult to know how seriously respondents take such questions and what degree of certainty their response implies (see van der Eijk et al., 2006 for a fuller discussion on opinion poll limitations).

Suppose a polling company could sample the entire population and obtain the

vote shares for each party. This would have a sample error of zero but would not be a precise forecast for the vote share at some future date, although presumably, it is an unbiased estimator due to the independence of potential news shocks. The nature of the polling business is to use this estimator as a proxy for future election day vote shares and extrapolate them into seat shares.




FIGURE 6.10: Opinion Poll Forecasts




FIGURE 6.11: Market Forecasts

Figures 6.10 and 6.11 indicates the consistency between the betting market and opinion polls estimation of party vote share predictions. The opinion polls have greater volatility than the mean implied vote shares from the betting market. Also, the opinion polls are just a point estimate (with a sampling error), whereas the betting market also produces a confidence band for these estimates vis the interval level bets.

Additionally, the betting markets ask the more interesting question with regards to pricing – what the is likely distribution of beliefs over the election date and takes into account the current voting intentions and the variance of changes that might occur due to exogenous events or last-minute switching. Consequently, although the mean of the betting market distributions is similar to the pollsters mean, the standard deviation is much higher – a standard deviation of approximately 6% versus the pollsters' typical reported (sample) error of 1.5%. As time progresses, the betting market distribution variance reduces as the possibility for exogenous shocks reduces, and this is indeed what the empirical data reports. The betting markets effectively ask the question: "how do you think other people will vote?" which is different from asking, "How would you vote?" On the day before the election, the 6 opinion polls were published, and the results are summarized in Table 6.4.

	Poll Average	Poll Range	Market %	Final %
Conservative	43.83%	41-45%	42.88%	43.60%
Labour	33.5%	31-35%	33.88%	32.2%
Lib Dem	11.33%	9-14%	13.13%	11.50%

TABLE 6.4: Opinion Poll Summary - Election Day -1

Although the average of the final six final polls published was very close to the final realization, the range of the polls' estimates was quite large – average absolute error of 4, which is the same absolute error as the betting market. The average RMSE of the pollsters' forecasts was 2.54% compared to the market RMSE of 2.44%, suggesting that the market forecast was slightly better than the average opinion poll. The betting market error could be an example of the bias inherent in prediction markets – the longshot bias of small probabilities, the home bias, which is the desire of gamblers to bet on their home outcome – seen in International football and US baseball. See [Wolfers and Zitzewitz \(2004\)](#) for a complete discussion of possible prediction market bias sources.

6.6 Model

6.6.1 Network Structure of the Betting Market

A network can be used to describe a market, where the nodes of the network represent the economic entities and the links represent potential trading relationships between them. In the gambling market, we have bookmakers and customers, and the customers connect to the bookmakers in some way and represent a possible trading relationship. This basic structure is known as a bipartite network (two node types), and we can calibrate this network with empirical data. The gambling market structure's main parameters can be inferred from the UK Gambling Commission's annual report and survey of gambling behavior. The Gambling Commission is the UK's main regulator for gambling activity and conducts regular surveys to track problem gambling and changes in gambling behavior. Since 2015, they have also tracked how gamblers hold many separate accounts, and in 2019 this was reported as 2.7 accounts per gambler, with 44% of gamblers having only one account. These are self-reported numbers

but appear to be steady over the last five years. (www.gamblingcommission.gov.uk). An unfortunate consequence of regulation to protect vulnerable gamblers is the more onerous account setup procedures which appear to have had the effect of slightly disincentivizing the opening of competing accounts – 3.5 to 2.7 over the five years. It should be noted that the majority of bookmakers for sports betting also make prices for political events.




FIGURE 6.12: Average Number of Online Accounts, taken from the Gambling Commission

The Gambling Commission reports that 44% of clients have only one account, with the mean number of accounts as 2.7. We can simplify this distribution as the clients belonging to 2 groups - a group that selects only one betting account and a second group that selects five accounts. That is, 44% of clients randomly choose only one bookmaker, and the second group of 56% of clients randomly chooses five bookmakers to produce a link distribution consistent with the Commission's data.

Clients choose their bookmakers for multiple reasons, including; tightness of spreads, range of betting products, online technology solutions, financial solvency, etc. The largest and most popular bookmakers spend over £1.5B ³ in the UK in advertising per year, and so it is reasonable to assume a certain amount of preferential selection for these bookmakers. i.e., the client who wishes to open one betting account is more likely to choose an already popular one. The combination of this client degree distribution (which is a common feature of many socio-economic networks) combined with the preferential selection for certain large bookmakers means that the largest bookmakers have a higher ratio of degree 1 clients to degree 5 clients than

³www.responsiblegamblingtrust.org.uk report on UK Gambling advertising

the peripheral bookmakers. This, in turn, means that the most popular bookmakers tend to price their odds towards the monopoly level price, and the less popular bookmakers, which are subject to more competition, tend to price their odds for a competitive market. Normally, when the price maker is more informed than the price taker, the monopoly price is higher than the competitive price. However, if the clients are equally or better informed than the bookmaker, the reverse can be true. This is due to the winners' curse dominating the payoffs. See section 6.9 for a complete discussion.

We will assume that at each time t , both bookmakers and clients have an equal standard deviation of beliefs in the final outcome - i.e. all players have the same degree of uncertainty as to the final value. We also make the assumption that there is preferential selection for the largest and most popular bookmakers and although we don't have the data for the numbers of clients, we use quoted UK gambling revenue figures as a proxy for popularity. Suppose bookmaker i has percentage of total market revenue of $RV_i = \text{UK gambling revenue} / \text{total UK gambling revenue}$, then the probability of bookmaker i being selected by a new client is equal to RV_i .

Suppose there are M clients and a fraction α clients are degree 1 and $(1 - \alpha)$ clients are degree 5, then each bookmaker B_i , r has an expected ratio f of 1 degree clients to 5 degree clients as

$$E[f] = \frac{\alpha RV_i}{(1 - \alpha)(1 - (1 - RV_i)^5)} \quad (6.1)$$

which is a strictly increasing function in α and also strictly increasing in RV_i , i.e. the ratio of 1-degree clients (no competition) to 5-degree clients (competition) increases with both the amount of preferential selection and with the proportion of 1-degree clients in the market and is the primary reason for centrality discounts and premiums in market networks

It follows that if the expected payoff for bookmaker i is:

$$E[\Pi] = M(\alpha RV_i E[\Pi_{deg1clients}] + (1 - \alpha)(1 - (1 - RV_i)^5) E[\Pi_{deg5clients}]) \quad (6.2)$$

Typically, the monopoly price is greater than the competitive price but in a common

value situation, the reverse can be true due to the winners' curse effect. If however, the optimal spread for 1-degree clients is lower than the optimal spread for 5-degree clients, the blended maximizing spread will be higher for less popular bookmakers and cause a centrality discount.

6.6.2 Multi-Period Unknown Common Value FPSBA Model

Suppose a bookmaker is connected to n customers, and each customer is connected to m ($n \gg m$) bookmakers. The bookmaker wants to make a betting price for a Conservative win in the General Election, defined as winning the most seats out of all the other parties, and the connected customer will place a unit bet with their connected bookmaker that has the best (lowest probability) price for this event. This is a Bertrand competition model in odds (probabilities) prices and can be analyzed using a common value FPSBA representation as described in the previous chapters.

The betting market most closely resembles the FPSBA as the bookmakers can only observe previous competitor prices (published historical odds) as there is no real-time centralized exchange to compare prices. Bookmakers do not observe the other (unknown) competitor prices during the actual odds quoting process, as the exact network structure is unknown, i.e., you do not know whom you are competing against in each quote. It is also a one-shot price – bookmakers cannot change their price while quoting with a customer and have no accurate knowledge of competitors' prices. A critical feature of ascending auctions is the ability to sequentially improve your price as other prices become apparent. For these reasons, we use an FPSBA as a modeling concept rather than an ascending auction process to model the trading interactions. However, it should be noted that the bookmakers do have more information than is usually assumed in a standard FPSBA due to the availability of historical quotes, which would allow them to infer likely competitor prices.

Although auction theory categorizes auctions as belonging to 4 main types (FPSB, Vickrey, Dutch, English), MaKafee et al. (1987), real-world auctions often have idiosyncrasies that deviate from the strict definitions. In the betting market, the two types that most closely resemble the pricing dynamics are an FPSBA and an English (ascending) auction. Although it appears possible for each bookmaker to observe the other bookmakers' prices (and hence adjust their own price/valuation) in

the manner of an ascending auction, because of the limited connectivity of the clients and the limited information of the network connections, bookmakers do not know which competitor prices are relevant in the competition.

The bookmaker forms beliefs over the likely win probability of each party, and we assume (see section 6.5 for the empirical evidence) that the vote share of each party is distributed as a normal random variable with the Conservative Vote Share, $C_{VS} \sim N(\mu_C, \sigma_C)$, Labour Vote Share, $L_{VS} \sim N(\mu_L, \sigma_L)$, which we assume are dependent normal random variables as a large reading for one will translate into a small reading for the other.

The probability p , that the Conservatives gain a higher vote share than Labour (discounting the LibDem and Brexit party chances) is given by

$$p = P(\text{Con wins}) = P(C_{VS} > L_{VS}) = 1 - \Phi \left(\frac{(\mu_L - \mu_C)}{\sqrt{\sigma_L^2 + \sigma_C^2 + 2\text{COV}(C_{VS}, L_{VS})}} \right) \quad (6.3)$$

Let $P_T \equiv \mu_C - \mu_L$ and $\sigma_D \equiv \sqrt{\sigma_L^2 + \sigma_C^2 + 2\text{cov}(C_{VS}, L_{VS})}$, then the distribution of the excess Conservative vote share is distributed as $C_{\text{excess}} \sim N(P_T, \sigma_D)$

Bookmaker i estimates this true value P_T , the excess vote share, on election day ($t=T$) with a certain error or confidence.

Since odds and probabilities are interchangeable, for simplicity we let the bookmaker make a probability price $p \in [0, 1]$ for the bet on the success of the Conservative party in the election.

As in previous chapters, we model this measurement or estimation process by bookmaker i , receiving a signal $S_{i,t} = P_T + \epsilon_t$, at time t , which is an estimate of the Conservative's excess vote share that realizes on election day with mean P_T , the true excess realized vote share. $\epsilon_t \sim N(0, \sigma_D)$ with variance σ_D

Bookmaker i , makes a betting price at time t of $B_{i,t} = S_{i,t} + \delta_i$ to his connected customers and these customers select the bookmaker with the lowest (probability) price to trade. δ_i is the overround or profit margin selected by the bookmaker for this bet.

We simplify the behaviour of customers as follows – for each customer of degree k , (so is connected to k bookmakers), each customer forms a belief R as to the correct

winning probability of the event. The customer then checks the prices of his connected bookmakers, and if the odds imply a lower probability of winning than their estimate R , a unit trade occurs.

The payoff for bookmaker i , in competition with $k-1$ other bookmakers j is:

$$\Pi_i = \begin{cases} B_{i,t} - P_T : \text{if } B_i < B_j \forall j \cap B_i < R \\ 0 : \text{otherwise} \end{cases} \quad (6.4)$$

i.e. bookmaker i makes a payoff of $B_{i,t} - P_T = \varepsilon_i + \delta_i$, if $B_{i,t}$ is the lowest price in the network of the customer connections and is also better than the customer's estimate R , which is a Bertrand price competition model in odds prices with a common unknown true value. This means that in equilibrium, the bookmaker with the highest adjusted signal error (the estimate plus the overround) is selected to trade.

In a market with a single client who receives a signal $S_i = P_T + \epsilon_i$ where ϵ_i is drawn from a normal distribution with zero mean and standard deviation σ_C and sets a reservation price equal to S_i (their estimate of the true probability) and n ($n \geq 1$) bookmakers, where the bookmakers receive an independent signal $S_i = P_T + \epsilon_i$ where the ϵ_i are drawn from a normal distribution with zero mean and SD σ_B , the expected payoff for a bookmaker i is a function of all the bookmakers' margins and the signal standard deviations of bookmakers and clients, moreover:

$$E[\Pi_i(\delta_i, \delta_{-i}, \sigma_B, \sigma_C)] = - \int_{-\infty}^{\infty} \frac{s}{\sigma_B} \left[\prod_{\substack{j=1 \\ j \neq i}}^n \Phi\left[\frac{s + \delta_j}{\sigma_B}\right] \right] \Phi\left[\frac{s}{\sigma_C}\right] \phi\left(\frac{s + \delta_i}{\sigma_B}\right) ds \quad (6.5)$$

Where δ_i is the bookmaker's profit margin, δ_{-i} are the profit margins of the competing bookmakers, σ_B is the standard deviation of the bookmakers' belief over the final point value estimate.

This model of expected payoffs, shown in figure 6.13, gives some insight into the reason the overround should never be zero and is an illustration of the winners' curse effect – the conditional probability of making a wrong (losing) price after having been selected to trade is greater than the unconditional probability of making a losing price.

The winner's curse can only be avoided by adding an overround to your estimate of true value and there exists an optimal amount of overround that needs to be added to




FIGURE 6.13: Illustration of Winner's Curse due to Overrounds

maximize expected payoffs. We assume that each bookmaker looks to maximize this expected payoff and are risk neutral, and hence only care about expected returns. This can be represented as a static Bayesian game and has a Bayesian Nash equilibrium solution as per the other chapter.

Summarizing the important assumptions of this model:

- Risk Neutral - bookmakers maximize expected payoffs and are indifferent to inventory and variance of payoffs
- Rationality - bookmakers make prices that are independently profitable – relative to their beliefs, each event odds bet has a positive expected payoff. This a consequence of 1 – the expected utility of a losing bet can be positive if it offsets some risk)
- All bookmakers have same variance of error in estimate of final values.
- Distribution of errors is common knowledge which is itself common knowledge.

At each time period, t , each bookmaker looks to add a profit margin $\delta_{i,t}$ to their expected true final outcome value and makes a price of $B_{i,t} = S_{i,t} + \delta_{i,t}$, and at each time period t , the mean and variance of the distribution of $S_{i,t}$ may change due to exogenous events. The excess margin $\delta_{i,t}$ need not be the same for each event and $\sum_E \delta_{i,t} = \pi > 1$ is summed over all events E and $\pi - 1$ is known as the overround.

Each bookmaker is simultaneously attempting to maximize their expected payoff

and theoretically there is an amount of margin that each bookmaker adds that cannot be unilaterally improved upon. This Bayesian Nash equilibrium (it is a Bayesian game because of the belief assumptions of all the players and the common prior assumption) can be solved numerically with the algorithm of the previous chapter.

In this competitive setup as described, at each day t , each bookmaker has a Nash Equilibrium amount, δ_i , to add to their estimated final value, and an amount $-\delta_i$ to subtract from it in order to generate a 2-way probability price in an event's outcome, given a common final event distribution function F .

The bid-ask spread is the difference between what prices a customer could back and lay an outcome with the same bookmaker. We investigate if the bookmakers' bid-ask spreads in odds prices approximate a Nash equilibrium bidding strategy by analyzing the empirical data and estimating the relevant variables required to calculate the theoretical spread levels. The expected payoff function includes three unknown quantities that we will calculate or estimate from the empirical dataset. Firstly, the functional form of the distribution of beliefs, which we shall approximate as being normal and is consistent with the empirical data, secondly, the mean and variance of these beliefs, which we also calculated from the data and lastly the number of competitors and network position each bookmaker faces when making a price to a customer. We infer this distribution from the UK gambling commission survey reports. We also take the UK betting revenue as a proxy for popularity (preferential selection) in constructing preferential attachment in the network to examine the centrality premium and behavioral differences.

6.6.3 Converting Vote Share to Seat Majority Beliefs

The aim now is to produce a seat win probability from the vote share data. As before, let conservative final vote share be distributed as $C \sim N(\mu_C, \sigma_C)$ and the expected Labour final vote share be $L \sim N(\mu_L, \sigma_L)$, then the probability p , that the Conservatives gain a higher vote share than Labour (discounting the LibDem chances) is given by

$$p = 1 - \Phi \left(\frac{(\mu_L - \mu_C)}{\sqrt{\sigma_L^2 + \sigma_C^2 + 2COV(C_{VS}, L_{VS})}} \right) \quad (6.6)$$

Unfortunately, the beliefs are not independent as a large realization for one party would likely produce a large negative realization for the other party as there is a fixed population to draw from, therefore the covariance between the Labour and Conservative vote shares, $\text{cov}(C,L)$, is unknown and certainly negative.

However, in order to calculate spreads, we just need the fact that the distribution is normal with some unknown variance σ_X since the distributions of both Labour and Conservative vote shares is normal, any linear combination of these distributions is also normal. The $p(C > L)$ is already known by the separately quoted win probabilities (point values) and μ_C and μ_L are known from the means of the interval level bet data (belief distributions). We can then solve for the effective variance of this new normal variable, which is $\sigma_X^2 = \sigma_L^2 + \sigma_C^2 + 2\text{COV}(C, L)$, to use for the effective variance of the outright win probabilities and the inputs to the theoretical spread calculations. This variable is the variance of the beliefs or margin of error of the win probabilities.




FIGURE 6.14: Standard Deviation of Conservative Party Most Seats Win Belief Probabilities

Figure 6.14 plots these standard deviations of Conservative most seat win probability beliefs that are derived from the Conservative and Labour vote share beliefs after solving for the implied covariance. It can be viewed as a confidence measure on the mean probabilities (estimates) of winning the most seats. Similar to financial

markets and standard behavioural theory, we notice over and under-reactions to exogenous shocks, which is particularly noticeable after the Brexit Party non-compete announcement on the 11th November. This was unequivocally good news for the Conservative vote share, which went up, but the initial surge in confidence appears exaggerated and was subsequently reversed over the next few days. The smoothed beliefs, also plotted, use 3-day simple moving average to give an indication of market beliefs that are more consistent with the observed changes in beliefs of the individual vote shares.

6.7 Observed and Predicted Spreads

We can now examine the predicted equilibriums in our betting market model and compare them with the data. We have now calculated or observed all of the parameters necessary for the bid-ask spread pricing model, namely, the extent of competition (network properties) and belief distributions (standard deviations). We now use these parameters to determine if the bookmaker bid-ask spreads in the market follow either a Nash equilibrium strategy or a zero profit strategy. For each day, we use the implied variances of winning beliefs (Conservatives only for brevity) to calculate both the NE spread and the zero-profit spread that are consistent with these parameters.




FIGURE 6.15: Mean Observed Spreads and Predicted Spread Levels Evaluated from Belief Data

Figure 6.15 shows that the average bookmaker spread tends to follow a strategy close to the theoretical NE spread calculated from the implied beliefs of winning. The predicted spreads are more volatile than the observed spreads in the market, and this may be due to data fluctuations in the belief odds data. In line with the theory on Bertrand competition on common value items, the bookmakers do not charge zero expected profit prices in equilibrium.

Figure 6.14 shows the level of noise in the beliefs and smoothed belief levels. Using the smoothed beliefs as inputs for the equilibrium level gives the following outputs for a predicted spread level versus the observed spread levels. Also plotted, for reference, is the calculated zero profit spread levels. These are calculated in the same way from the payoff functions but with the algorithm target set to zero profits.

Looking instead at the spreads derived from smoothing the belief data, Figure 6.16 gives a very clear result – the bookmakers' spread strategy appears to follow these predicted NE spreads quite closely and are much higher than the estimated zero-profit spread levels. The smoothing uses a simple 3-day moving average of the beliefs.




FIGURE 6.16: Observed and Belief Smoothed Predicted Spread Levels

6.8 Model Evaluation

A common and straightforward approach to evaluate models on complex systems is to regress predicted vs observed values, Piñeiro et al. (2008) (or vice versa) and compare slope and intercept parameters against the 1:1 line. We ran a linear regression between the observed spread prices and the predicted equilibrium spreads using the market beliefs of the final outcome variance as inputs. The model produced a nearly 1:1 relationship with an R^2 of 90%. A perfect model would have a slope of 1 and an intercept of zero, and the fitted linear model has a slope of 1.118 with 95% confidence bands of (0.9036, 1.332) and an intercept of -0.5888 with confidence bands of (-1.951, 0.7729).




FIGURE 6.17: Observed vs Predicted Linear Regression

Figure 6.17 plots the linear model of the predicted versus the observed values.




FIGURE 6.18: Residuals of observed vs predicted values regression model

The plot of the residuals, show in figure 6.18, also shows no obvious structure. Conducting a 'runs' test, Bradley (1968), on the residuals, where the null hypothesis H_0 : the sequence was produced in a random manner vs H_1 : the sequence was not

produced in a random manner was investigated. The result is we cannot reject the null hypothesis at both the 95 and 99% level.

R^2 also represents the proportion of the linear covariance of observed and predicted values with respect to the total variance of observed and predicted values. In this sense, the R^2 indicates how much of the linear variation of observed values is explained by the variation of predicted values. Linearity between observed and predicted values can be tested following (Smith and Rose, 1995). Thus, the R^2 of observed vs. predicted values is a valid parameter that gives important information of the model performance

It is always possible that our “explanatory variables” are completely useless for predicting the observed values and are due to natural variability of the data - noise vs signal. We can formulate this as a hypothesis test with the null hypothesis that all regression parameters (except the intercept) are zero, that is, we form the null hypothesis:

$$H_0 : \beta_1 = 0$$

versus the alternative hypothesis:

$$H_A : \beta_1 \neq 0$$

That is, we can say that under H_A , the model has some predictive power over a constant model. This is accomplished by performing an F-test on the regression coefficients of the predicted vs observed values.

Linear regression model:

$$y \sim 1 + x1$$

Estimated Coefficients:

	Estimate	SE	tStat	pValue
(Intercept)	-0.58883	0.63032	-0.93418	0.36724
x1	1.1176	0.099093	11.279	4.3792e-08

Number of observations: 15, Error degrees of freedom: 13

Root Mean Squared Error: 0.391

R-squared: 0.907, Adjusted R-Squared: 0.9

F-statistic vs. constant model: 127, p-value = 4.38e-08

FIGURE 6.19: F-test results on the linear regression coefficients

The F-statistic vs a constant model of 127 says we can reject the null hypothesis H_0 at both the 95% and 99% confidence level and has a p-value of 4×10^{-8} . That is, the probability of a type 1 error (data is actually generated from a constant model) is 4×10^{-8} . Clearly as we have only one explanatory variable, the predicted value, the p-value of the F-statistic is the same as the t-statistic.

Given the assumptions, noisy data and complexities of calculations, it is notable that the price competition mechanism over a network manages to produce prices that are measurably close to the theoretical predictions.

6.9 Centrality Premium Effect

In many financial markets, a centrality premium exists; that is, the more centrally located dealers charge higher prices than the more peripheral dealers. The betting market structure is very similar to the OTC market structures of the financial markets, and we examine the data to determine if any centrality discount or premium exists. The results show a centrality discount for the more central and active bookmakers, compared to the smaller, less active ones. This discount is predicted in our common value network model when there are symmetric information sets and preferential selection by the customers for a particular subset of bookmakers. It seems reasonable to assume that given the costs of establishing accounts and the average number of accounts being three, that the biggest bookmakers are likely to be preferred as one of the client choices over the smaller ones when opening only a small number of accounts. This preferential selection is the driving force of the centrality discount, as presented in the previous chapter.

The 16 active bookmakers that were represented on oddschecker.com (3 bookmakers making sporadic prices were dropped from the analysis) were sorted by total UK revenue using published online data.

Table in figure 6.20 shows the ranking and average spreads over the 1-month period prior to the election day.

As time progresses to the election day, the outcome's variance reduces due to the reduced probability of election changing news – leader's debates, speech gaffes, for example. Since our model estimates a spread that is positively related to the final

Company	Rev (£M)	Av Spread	Company	Rev (£M)	Av Spread
bet365	2981	6.11%	Paddypower	533	6.26%
William Hill	1581	7.15%	sportingbet	400	6.16%
Coral	1507	4.88%	Boyle sports	300	7.44%
Ladbrokes	1507	4.88%	Betfair	393	6.10%
Betvictor	1000	6.90%	Betway	173	8.09%
Unibet	751	4.83%	Marathon	55	8.55%
Betfred	728	4.29%	SpreadEX	50	6.59%
sky bet	670	6.06%	Royal Panda	25	6.91%
Average	1341	5.64%	Average	241	7.01%

FIGURE 6.20: Partitioning on Bookmakers by Revenue

value estimate variance, we would expect a drift down in spreads on no significant news and a sharp drop in spreads on any large positive news that permanently increases a party's chance of election. The data suggest that this effect does occur – there is a significantly large drop in spreads after the Brexit party announcement on the 11th November that they were not going to compete in Conservative-held seats. This produced a significant increase in the expected Conservative Vote share and a considerable reduction in a Conservative win's uncertainty. Spreads for a conservative win collapsed by nearly 2% - the most significant drop over the whole election dataset. It could be argued from the data that the Conservatives effectively won the election on 11th November with the Brexit party non-compete announcement, as demonstrated by the large reduction in the belief variances

Figure 6.21 plots the average bid-ask spreads of both central and peripheral bookmakers and clearly shows that the central bookmakers consistently make tighter spreads than the peripheral bookmakers. This is a surprising result, considering the extra costs of marketing that the large bookmakers spend. This result, however, is consistent with our model of spread prices with a client-base that is at least as informed as to the outcome as the bookmakers and bookmakers are subject to a level of preferential selection as new clients enter the market. In a growing network with preferential selection, the bookmakers with the largest numbers of links are more likely to get selected by the clients. Following the network model described in section 6.6.1, we have that each client has a 44% probability of having just one link and a 56% probability of having five links. The bookmakers are selected with some preferential




FIGURE 6.21: Spreads of Largest and Smallest Bookmakers

selection that is based on their popularity. At each time point, we assume that the clients and the bookmakers have an equal standard deviation of beliefs as to the true final value.

The bookmakers all seek to maximize their expected profits by adjusting their profit margin δ , which is given by:

$$E[\Pi(\delta)] = M(0.44RV_i E[\Pi_{deg1clients}(\delta)] + (0.56)(1 - (1 - RV_i)^5) E[\Pi_{deg5clients}(\delta)]) \quad (6.7)$$

where M is the total number of clients, RV_i is the proportion of market share of bookmaker i .

and the expected payoff from degree 1 clients of bookmaker i , is given by:

$$E[\Pi_{i,deg1clients}, (\delta_i, \sigma_B)] = - \int_{-\infty}^{\infty} \frac{t}{\sigma_B} \Phi\left[\frac{t}{\sigma_B}\right] \phi\left(\frac{t + \delta_i}{\sigma_B}\right) dt \quad (6.8)$$

and the expected payoff from degree 5 clients is :

$$E[\Pi_i(\delta_i, \delta_{-i}, \sigma_B)] = - \int_{-\infty}^{\infty} \frac{t}{\sigma_B} \left[\prod_{\substack{j=1 \\ j \neq i}}^5 \Phi\left[\frac{t + \delta_j}{\sigma_B}\right] \right] \Phi\left[\frac{t}{\sigma_B}\right] \phi\left(\frac{t + \delta_i}{\sigma_B}\right) dt \quad (6.9)$$

These systems of equations (over all bookmakers) have a fixed point at which each

bookmaker is in equilibrium regarding the profit margin added to their estimate and is a Bayesian Nash equilibrium that can be approximated numerically.

If bookmakers and clients are, on average, equally able to measure the outcome, the equilibrium spread level for the 1-degree clients is lower than the equilibrium spread level for the 5-degree clients and was discussed at length in Chapter 3. This common variance tends to cause a centrality discount in equilibrium when there is a preferential selection for the bookmakers.

6.10 Conclusion

The UK general election provided a perfect dataset to test the idea in this thesis that non-centralized markets with common uncertainty over the true value can produce non-intuitive pricing effects in equilibrium. We collected data over the month preceding the election and showed that beliefs over the final outcome were well represented at each time point by a normal distribution, which gives a richer representation of beliefs than the usual point estimates. We then used these belief distributions to create a Bertrand unknown common value competition model in odds prices and found that the quoted prices accorded very well with the model's theoretical Nash equilibrium levels. Similar to some OTC financial markets (the securitization market analyzed by [Hollifield et al. \(2012\)](#) also exhibited a centrality discount), we also found a centrality discount in the prices charged by the most popular (central) bookmakers in the gambling market. Although the centrality discount effect is not totally explained using traditional economic theory, our model offers a possible explanation - it is simply a property of a Bertrand competition model when played over a network with preferential attachment.

In terms of forecasting and informational content, we find, similar to others, that the betting markets provide a forecast at least as accurate as polling data. One caveat is that there is a well-documented longshot bias for small probability events that need to be considered, and this was adjusted for by standard normalization methods.

Chapter 7

Conclusion and Discussion

This thesis describes a simple network model of the trading process in a general client-dealer OTC market, with a common unknown asset value. Clients connect to dealers via a network of trading relationships, and both clients and dealers form an estimate of this asset's true value. Clients request quotes from their connected dealers, and dealers respond with a firm quote to buy and sell, and the client order transacts at the best-observed quoted price.

Our model offers insight into the interesting equilibrium pricing between perfect competition and monopoly by examining the network topology within a framework of asymmetric information sets. We view dealer bid-ask spreads as analogous to the signal reductions in first-price auctions and use traditional auction analysis to find the optimal dealer pricing strategy. We jointly maximize the dealer payoff functions numerically to find their Bayesian Nash equilibrium strategy and examine the effects of different network configurations and asymmetric information on the equilibrium bid-ask spread.

As in many traditional asymmetric information microstructure models, we find asymmetric information to be a significant driving force in determining equilibrium dealer spreads. The extent of asymmetric information is a primary driver of why in a network with heterogeneous client links, central dealer spreads can be either a premium or a discount to the less-central dealer spreads. Setting bid-ask spreads to either compensate or exploit informational asymmetries depends on both the relative information sets and network topology, which interact in non-linear ways.

Similarly, we find that the client degree distribution is a major factor in determining the effects of increasing information transparency on equilibrium bid-ask spreads,

where transparency is modeled as a refining of the clients' information sets regarding the true value of the asset. Clients with low degree (less than 3) experience reduced spreads, whereas higher degree clients face higher spreads when transparency is increased. This is due in part to the increasing dominance of the winners' curse effect with improved client information sets and dealers attempting to mitigate this in equilibrium by increasing spreads. This may offer some explanation as to the mixed empirical results seen when market transparency has been increased by regulators and is further evidence of the importance that market network topology plays.

This common value model also generates two types of price dispersion in both the buy and sell prices. Firstly, from the asset value estimate variance, with increased dispersion when asset value uncertainty is high. Price dispersion is also shown to fall when clients become better informed as to the true value of the asset. This price dispersion is a feature unique to non-centralized markets, with different prices trading in different parts of the network. Secondly, there is also a prices dispersion between clients of different degree, with increased degree equating to lower price dispersion as the variance of the equilibrium prices decreases with degree.

We use the equilibrium results to show how our model can explain some other puzzles in the empirical literature inconsistent with established economic theory. We looked at how spreads in low-risk municipal bonds can be greater than medium-risk bonds, which can be greater still than risky equities. Additionally, illiquid bonds can have tighter spreads than liquid bonds. These phenomena cannot be explained by pure informational asymmetry or inventory arguments. In harder-to-value illiquid bonds, inventory models would predict the opposite effect due to the increased difficulty in offsetting positions and so would command a higher spread. Similarly, asymmetric information models predict that the harder to value securities would command a higher bid-ask spread. When the network connections are included, the combination of clients' low-degree and relative uninformedness explains these observations.

Finally, we apply the network model to the gambling market with data gathered from the UK general election of 2019. We find that a normal distribution well represents the distribution of beliefs of the true value, and the network topology is similar to many OTC financial markets. We demonstrate that bid-ask spreads in odds prices

are consistent with our network model.

We have shown that even a simple Bertrand price competition can give rise to many non-linear and counter-intuitive effects when executed over a network. The three elements of this story, information sets, degree distribution, and preferential selection, combine to explain empirical data that standard linear models do not easily describe, and our network model describes a parsimonious model of trading in OTC markets.

Real markets, however, are generally not precisely represented by one-shot games but by multi-period transactions, where future relationships and learning become an essential consideration. The network structure in this model does not change - clients still prefer specific dealers and have a small number of links that drive the topology, but clients and dealers can learn from the traded prices and possibly change their network links.

The standard model of rational learning in a network maintains that individuals use Bayes' rule to incorporate any new piece of information into their beliefs (Molavi, P., Tahbaz-Salehi, A. and Jadbabaie, A. (2018)) and involves each agent forming a prior belief updated by information on each trade observed. Even in simple networks, this is an onerous task as it requires updating beliefs on all other agents' information sets in the network at every time period. Other heuristic methods of learning have been devised and the leading behavioral model here is the DeGroot heuristic model, deGroot (1974), which is a simple method whose results have been empirically tested and shown to outperform Bayesian learning in experimental work, (Chandrasekhar, A.G., Larreguy, H. and Xandri, J.P. (2020)). This method involves simple averaging of one's own and one's neighbors' beliefs as new data is revealed. Similarly, the dealers' strategic nature means that the multi-period game alters the pricing strategy, since the dealers' objective function is now the sum of transactions over a long period. For example, a dealer may price below equilibrium levels to win new customers (network links) in a dynamic network, drive out competitors, and exploit their superior network position at a later stage of the game. This predatory pricing phenomena was examined by Milgrom (1991). In addition, any agent with a superior estimate of the true value might disguise this fact to retain a competitive advantage because of learning by other agents.

Other avenues for future work include applying this model to the other types of bipartite market networks, such as the client-client network, where boundedly rational clients trade with other clients without intermediation, similar to an eBay or other peer-to-peer network type structure. In the dealer-dealer networks, this is more complicated as the strategic dealers fall foul of various no-trade theorems, for example, the seminal no-trade theorem of Myerson, Roger B.; Mark A. Satterthwaite (1983). The common prior assumption ensures that in equilibrium, entirely strategic dealers with knowledge of the probability distribution of signals will have no pure financial incentive to trade with each other, (Milgrom, Paul; Stokey, Nancy (1982)), so an inventory management element would need to be introduced in order to incentivize trade.

Bibliography

- J. Abad, I. Aldasoro, C. Aymanns, M. D’Errico, L. Fache Rousová, P. Hoffmann, S. Langfield, M. Neychev, and T. Roukny. Shedding light on dark markets: First insights from the new eu-wide otc derivatives dataset. (11), 2016. URL <https://EconPapers.repec.org/RePEc:srk:srkops:201611>.
- G. Afonso and R. Lagos. Trade dynamics in the market for federal funds. *Econometrica*, 83(1):263–313, 2015. doi: <https://doi.org/10.3982/ECTA10586>. URL <https://onlinelibrary.wiley.com/doi/abs/10.3982/ECTA10586>.
- R. Albert and A.-L. Barabási. Statistical mechanics of complex networks. *Reviews of Modern Physics*, 74(1):47–97, Jan 2002. ISSN 1539-0756. doi: 10.1103/revmodphys.74.47. URL <http://dx.doi.org/10.1103/RevModPhys.74.47>.
- Y. Amihud and H. Mendelson. Dealership market: Market-making with inventory. *Journal of Financial Economics*, 8(1):31–53, 1980. URL <https://EconPapers.repec.org/RePEc:eee:jfinec:v:8:y:1980:i:1:p:31-53>.
- K. J. Arrow, R. Forsythe, M. Gorham, R. Hahn, R. Hanson, J. O. Ledyard, S. Levmore, R. Litan, P. Milgrom, F. D. Nelson, G. R. Neumann, M. Ottaviani, T. C. Schelling, R. J. Shiller, V. L. Smith, E. Snowberg, C. R. Sunstein, P. C. Tetlock, P. E. Tetlock, H. R. Varian, J. Wolfers, and E. Zitzewitz. The promise of prediction markets. *Science*, 320(5878):877–878, 2008. ISSN 0036-8075. doi: 10.1126/science.1157679. URL <https://science.sciencemag.org/content/320/5878/877>.
- A. Babus and P. Kondor. Trading and information diffusion in over-the-counter markets. *Econometrica*, 86(5):1727–1769, 2018. doi: <https://doi.org/10.3982/ECTA12043>. URL <https://onlinelibrary.wiley.com/doi/abs/10.3982/ECTA12043>.

- J. Berg, R. Forsythe, F. Nelson, and T. Rietz. Results from a dozen years of election futures markets research. volume 1, Part 5, chapter 80, pages 742–751. Elsevier, 1 edition, 2008. URL <https://EconPapers.repec.org/RePEc:eee:expchp:5-80>.
- D. Bernhardt and E. Hughson. Splitting orders. Working Paper 888, Economics Department, Queen’s University, 1993. URL <https://EconPapers.repec.org/RePEc:qed:wpaper:888>.
- J. Bertrand. Book review of *theorie mathematique de la richesse sociale* and of *recherches sur les principes mathematiques de la theorie des richesses*. *Journal de Savants*, (67):499–508, 1883.
- H. Bessembinder, C. Spatt, and K. Venkataraman. A survey of the microstructure of fixed-income markets. *Journal of Financial and Quantitative Analysis*, 55(1):1–45, 2020. doi: 10.1017/S0022109019000231.
- B. Biais, L. Glosten, and C. Spatt. Market microstructure: A survey of microfoundations, empirical results, and policy implications. *Journal of Financial Markets*, 8(2):217–264, 2005. URL <https://EconPapers.repec.org/RePEc:eee:finmar:v:8:y:2005:i:2:p:217-264>.
- B. Biais, F. Declerck, J. Dow, R. Portes, and E.-L. von Thadden. European corporate bond markets: Transparency, liquidity, efficiency. *Centre for Economic Policy Research*, may 2006.
- E. Boehmer, G. Saar, and L. YU. Lifting the veil: An analysis of pre-trade transparency at the nyse. *Journal of Finance*, 60:783–815, 02 2005. doi: 10.2139/ssrn.368102.
- N. P. B. Bollen, T. Smith, and R. E. Whaley. Modeling the bid/ask spread: measuring the inventory-holding premium. *Journal of Financial Economics*, 72(1):97–141, 2004. URL <https://EconPapers.repec.org/RePEc:eee:jfinec:v:72:y:2004:i:1:p:97-141>.
- P. Bonacich. Power and centrality: A family of measures. *American Journal of Sociology*, 92(5):1170–1182, 1987. ISSN 00029602, 15375390. URL <http://www.jstor.org/stable/2780000>.

- S. P. Borgatti and M. G. Everett. Models of core/periphery structures. *Social Networks*, 21(4):375–395, 2000. ISSN 0378-8733. doi: [https://doi.org/10.1016/S0378-8733\(99\)00019-2](https://doi.org/10.1016/S0378-8733(99)00019-2). URL <https://www.sciencedirect.com/science/article/pii/S0378873399000192>.
- M. W. Brandt and K. A. Kavajecz. Price discovery in the u.s. treasury market: The impact of orderflow and liquidity on the yield curve. *The Journal of Finance*, 59(6):2623–2654, 2004. ISSN 00221082, 15406261. URL <http://www.jstor.org/stable/3694784>.
- K. Burdett and K. L. Judd. Equilibrium price dispersion. *Econometrica*, 51(4):955–969, 1983. ISSN 00129682, 14680262. URL <http://www.jstor.org/stable/1912045>.
- E. B. Cammack. Evidence on bidding strategies and the information in treasury bill auctions. *Journal of Political Economy*, 99(1):100–130, 1991. ISSN 00223808, 1537534X. URL <http://www.jstor.org/stable/2937714>.
- S. Clarke. Adjusting bookmaker’s odds to allow for overround. *American Journal of Sports Science*, 5:45, 01 2017. doi: 10.11648/j.ajss.20170506.12.
- C. Cooper and A. Frieze. A general model of web graphs. *Random Structures & Algorithms*, 22(3):311–335, 2003. doi: <https://doi.org/10.1002/rsa.10084>. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/rsa.10084>.
- T. E. Copeland and D. Galai. Information effects on the bid-ask spread. *The Journal of Finance*, 38(5):1457–1469, 1983. doi: <https://doi.org/10.1111/j.1540-6261.1983.tb03834.x>. URL <https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1540-6261.1983.tb03834.x>.
- B. Craig and G. von Peter. Interbank tiering and money center banks. *Journal of Financial Intermediation*, 23(3):322–347, 2014. URL <https://EconPapers.repec.org/RePEc:eee:jfinin:v:23:y:2014:i:3:p:322-347>.
- M. de Kamps, D. Ladley, and A. Simaitis. Heterogeneous Beliefs in Over-The-Counter Markets. (13/03), August 2012. URL <https://ideas.repec.org/p/lec/leecon/13-03.html>.

- M. H. DeGroot. Reaching a consensus. *Journal of the American Statistical Association*, 69(345):118–121, 1974. ISSN 01621459. URL <http://www.jstor.org/stable/2285509>.
- M. Di Maggio, A. Kermani, and Z. Song. The value of trading relations in turbulent times. *Journal of Financial Economics*, 124(2):266–284, 2017. URL <https://EconPapers.repec.org/RePEc:eee:jfinec:v:124:y:2017:i:2:p:266-284>.
- D. Duffie, N. Garleanu, and L. H. Pedersen. Over-the-Counter Markets. *Econometrica*, 73(6):1815–1847, November 2005. URL <https://ideas.repec.org/a/ecm/emetrp/v73y2005i6p1815-1847.html>.
- D. Duffie, P. DWORCZAK, and H. ZHU. Benchmarks in search markets. *The Journal of Finance*, 72(5):1983–2044, 2017. doi: <https://doi.org/10.1111/jofi.12525>. URL <https://onlinelibrary.wiley.com/doi/abs/10.1111/jofi.12525>.
- D. Easley and M. O’Hara. Price, trade size, and information in securities markets. *Journal of Financial Economics*, 19(1):69–90, 1987. URL <https://EconPapers.repec.org/RePEc:eee:jfinec:v:19:y:1987:i:1:p:69-90>.
- A. K. Edwards, L. E. Harris, and M. S. Piwowar. Corporate bond market transaction costs and transparency. *The Journal of Finance*, 62(3):1421–1451, 2007. ISSN 00221082, 15406261. URL <http://www.jstor.org/stable/4622305>.
- P. Erdős and A. Rényi. On random graphs i. *Publicationes Mathematicae Debrecen*, 6:290, 1959.
- E. Estrada and P. Knight. *A First Course in Network Theory*. Oxford University Press, United Kingdom, Mar. 2015. ISBN 978-0-19-872645-6.
- E. Franck, E. Verbeek, and S. Nüesch. Prediction accuracy of different market structures — bookmakers versus a betting exchange. *International Journal of Forecasting*, 26:448–459, 11 2009. doi: 10.1016/j.ijforecast.2010.01.004.
- D. M. Gale and S. Kariv. Financial networks. *American Economic Review*, 97(2): 99–103, May 2007. doi: 10.1257/aer.97.2.99. URL <https://www.aeaweb.org/articles?id=10.1257/aer.97.2.99>.

- M. B. Garman. Market microstructure. *Journal of Financial Economics*, 3(3): 257–275, 1976. URL <https://EconPapers.repec.org/RePEc:eee:jfinec:v:3:y:1976:i:3:p:257-275>.
- L. R. Glosten and P. R. Milgrom. Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. *Journal of Financial Economics*, 14(1):71–100, 1985. ISSN 0304-405X. doi: [https://doi.org/10.1016/0304-405X\(85\)90044-3](https://doi.org/10.1016/0304-405X(85)90044-3). URL <https://www.sciencedirect.com/science/article/pii/0304405X85900443>.
- D. K. Gode and S. Sunder. Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality. *Journal of Political Economy*, 101(1):119–137, 1993. ISSN 00223808, 1537534X. URL <http://www.jstor.org/stable/2138676>.
- M. A. Goldstein, E. S. Hotchkiss, and E. R. Sirri. Transparency and Liquidity: A Controlled Experiment on Corporate Bonds. *Review of Financial Studies*, 20(2):235–273, 2007. URL <https://ideas.repec.org/a/oup/rfinst/v20y2007i2p235-273..html>.
- J. Harsanyi. Games with incomplete information played by "bayesian" players, i-iii part i. the basic model. *Management Science*, 14(3):159–182, 1967. URL <https://EconPapers.repec.org/RePEc:inm:ormnsc:v:14:y:1967:i:3:p:159-182>.
- T. Hendershott, D. Li, D. Livdan, and N. Schürhoff. Relationship Trading in Over-the-Counter Markets. *Journal of Finance*, 75(2):683–734, April 2020. doi: 10.1111/jofi.12864. URL <https://ideas.repec.org/a/bla/jfinan/v75y2020i2p683-734.html>.
- T. Ho and H. Stoll. Optimal dealer pricing under transactions and return uncertainty. *Journal of Financial Economics*, 9(1):47–73, 1981. URL <https://EconPapers.repec.org/RePEc:eee:jfinec:v:9:y:1981:i:1:p:47-73>.
- T. S. Y. Ho and H. Stoll. The dynamics of dealer markets under competition. *Journal of Finance*, 38(4):1053–74, 1983. URL <https://EconPapers.repec.org/RePEc:bla:jfinan:v:38:y:1983:i:4:p:1053-74>.

- B. Hollifield, A. Neklyudov, and C. Spatt. Bid-ask spreads and the pricing of securitizations: 144a vs. registered securitizations. *SSRN Electronic Journal*, 03 2012. doi: 10.2139/ssrn.2023467.
- R. D. Huang and R. W. Masulis. FX Spreads and Dealer Competition across the 24-Hour Trading Day. *Review of Financial Studies*, 12(1):61–93, 1999. URL <https://ideas.repec.org/a/oup/rfinst/v12y1999i1p61-93.html>.
- M. O. Jackson and B. W. Rogers. Meeting strangers and friends of friends: How random are social networks? *American Economic Review*, 97(3):890–915, June 2007. doi: 10.1257/aer.97.3.890. URL <https://www.aeaweb.org/articles?id=10.1257/aer.97.3.890>.
- M. Janssen and E. Rasmusen. Bertrand Competition Under Uncertainty. (CIRJE-F-117), June 2001. URL <https://ideas.repec.org/p/tky/fseries/2001cf117.html>.
- J. H. Kagel and D. Levin. The winner’s curse and public information in common value auctions. *The American Economic Review*, 76(5):894–920, 1986. ISSN 00028282. URL <http://www.jstor.org/stable/1816459>.
- D. Kahneman and A. Tversky. Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2):263–291, 1979. ISSN 00129682, 14680262. URL <http://www.jstor.org/stable/1914185>.
- L. Katz. A new status index derived from sociometric analysis. *Psychometrika*, 18(1):39–43, 1953. URL <https://EconPapers.repec.org/RePEc:spr:psycho:v:18:y:1953:i:1:p:39-43>.
- A. Kolotilin and V. Panchenko. Estimation of a Scale-Free Network Formation Model. (2018-10), June 2018. URL <https://ideas.repec.org/p/swe/wpaper/2018-10.html>.
- R. E. Kranton and D. F. Minehart. A theory of buyer-seller networks. *American Economic Review*, 91(3):485–508, June 2001. doi: 10.1257/aer.91.3.485. URL <https://www.aeaweb.org/articles?id=10.1257/aer.91.3.485>.

- J. B. Krawczyk and S. Uryasev. Relaxation algorithms to find nash equilibria with economic applications. *Environmental Modeling & Assessment*, 5(1):63–73, 2000. doi: 10.1023/A:1019097208499. URL <https://doi.org/10.1023/A:1019097208499>.
- A. S. Kyle. Continuous auctions and insider trading. *Econometrica*, 53(6):1315–1335, 1985. ISSN 00129682, 14680262. URL <http://www.jstor.org/stable/1913210>.
- R. Lagos and G. Rocheteau. Liquidity in asset markets with search frictions. *Econometrica*, 77(2):403–426, 2009. ISSN 00129682, 14680262. URL <http://www.jstor.org/stable/40263870>.
- M. J. Lebo and H. Norpoth. The pm and the pendulum: Dynamic forecasting of british elections. *British Journal of Political Science*, 37(1):71–87, 2007. ISSN 00071234, 14692112. URL <http://www.jstor.org/stable/4497280>.
- D. Levin and J. Smith. Some evidence on the winner’s curse: Comment. *American Economic Review*, 81(1):370–75, 1991. URL <https://EconPapers.repec.org/RePEc:aea:aecrev:v:81:y:1991:i:1:p:370-75>.
- S. Levitt. Why are gambling markets organized so differently from financial markets? *Economic Journal*, 114:223–246, 02 2004. doi: 10.1111/j.1468-0297.2004.00207.x.
- D. Li and N. Schürhoff. Dealer Networks. *Journal of Finance*, 74(1):91–144, February 2019. doi: 10.1111/jofi.12728. URL <https://ideas.repec.org/a/bla/jfinan/v74y2019i1p91-144.html>.
- A. Madhavan, D. Porter, and D. Weaver. Should securities markets be transparent? *Journal of Financial Markets*, 8:265–287, 02 2005. doi: 10.1016/j.finmar.2005.05.001.
- D. Mallaburn, M. Roberts-Sklar, and L. Silvestri. Resilience of trading networks: evidence from the sterling corporate bond market. (813), August 2019. URL <https://ideas.repec.org/p/boe/boeewp/0813.html>.
- P. R. Milgrom and R. J. Weber. A theory of auctions and competitive bidding. *Econometrica*, 50(5):1089–1122, 1982. ISSN 00129682, 14680262. URL <http://www.jstor.org/stable/1911865>.

- S. Nadarajah and S. Kotz. Exact distribution of the max/min of two gaussian random variables. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, 16(2):210–212, 2008. doi: 10.1109/TVLSI.2007.912191.
- H. Nikaidô and K. Isoda. Note on non-cooperative convex game. *Pacific Journal of Mathematics*, 5:807–815, 1955.
- H. J. Paarsch. Deciding between the common and private value paradigms in empirical models of auctions. *Journal of Econometrics*, 51(1):191–215, 1992. ISSN 0304-4076. doi: [https://doi.org/10.1016/0304-4076\(92\)90035-P](https://doi.org/10.1016/0304-4076(92)90035-P). URL <https://www.sciencedirect.com/science/article/pii/030440769290035P>.
- L. Pedersen, D. Duffie, and N. Garleanu. Valuation in over-the-counter markets. *Review of Financial Studies*, 20:1865–1900, 02 2007. doi: 10.1093/rfs/hhm037.
- D. Peel, M. Cain, and D. Law. The favourite-longshot bias, bookmaker margins and insider trading in a variety of betting markets. *Bulletin of Economic Research*, 55:263–273, 02 2003. doi: 10.1111/1467-8586.00174.
- G. Piñeiro, S. Perelman, J. P. Guerschman, and J. M. Paruelo. How to evaluate models: Observed vs. predicted or predicted vs. observed? *Ecological Modelling*, 216(3):316–322, 2008. ISSN 0304-3800. doi: <https://doi.org/10.1016/j.ecolmodel.2008.05.006>. URL <https://www.sciencedirect.com/science/article/pii/S0304380008002305>.
- M. Rosenbaum. Betting and the 1997 british general election. *Politics*, 19(1):9–14, 1999. doi: 10.1111/1467-9256.00080. URL <https://doi.org/10.1111/1467-9256.00080>.
- D. Rothschild. Forecasting elections: Comparing prediction markets, polls, and their biases. *The Public Opinion Quarterly*, 73(5):895–916, 2009. ISSN 0033362X, 15375331. URL <http://www.jstor.org/stable/40467652>.
- R. R. Routledge and R. A. Edwards. Ambiguity and price competition. *Theory and Decision*, 88(2):231–256, 2020. URL https://EconPapers.repec.org/RePEc:kap:theord:v:88:y:2020:i:2:d:10.1007_s11238-019-09725-4.

- E. Snowberg and J. Wolfers. Explaining the favorite–long shot bias: Is it risk-love or misperceptions? *Journal of Political Economy*, 118(4):723–746, 2010. ISSN 00223808, 1537534X. URL <http://www.jstor.org/stable/10.1086/655844>.
- D. F. Spulber. Bertrand competition when rivals’ costs are unknown. *The Journal of Industrial Economics*, 43(1):1–11, 1995. ISSN 00221821, 14676451. URL <http://www.jstor.org/stable/2950422>.
- M. P. H. Stumpf and M. A. Porter. Critical truths about power laws. *Science*, 335(6069):665–666, 2012. ISSN 0036-8075. doi: 10.1126/science.1216142. URL <https://science.sciencemag.org/content/335/6069/665>.
- J. Surowiecki. The wisdom of crowds: Why the many are smarter than the few. london. *Abacus*, page 295, 01 2004.
- S. E. Thiel. Some evidence on the winner’s curse. *The American Economic Review*, 78(5):884–895, 1988. ISSN 00028282. URL <http://www.jstor.org/stable/1807155>.
- T. Valente, K. Coronges, C. Lakon, and E. Costenbader. How correlated are network centrality measures? *Connections (Toronto, Ont.)*, 28:16–26, 01 2008.
- V. Vovk and F. Zhdanov. Prediction with expert advice for the brier game. *Journal of Machine Learning Research*, 10, 10 2007. doi: 10.1145/1390156.1390295.
- M. Wall, L. Sudulich, and K. Cunningham. What are the odds? using constituency-level betting markets to forecast seat shares in the 2010 uk general elections. *Journal of Elections, Public Opinion and Parties*:3–26, 02 2012. doi: 10.1080/17457289.2011.629727.
- A. Wambach. Bertrand competition under cost uncertainty. *International Journal of Industrial Organization*, 17(7):941–951, 1999. ISSN 0167-7187. doi: [https://doi.org/10.1016/S0167-7187\(98\)00008-3](https://doi.org/10.1016/S0167-7187(98)00008-3). URL <https://www.sciencedirect.com/science/article/pii/S0167718798000083>.
- X. Wang and M. Pleimling. Online gambling of pure chance: Wager distribution, risk attitude, and anomalous diffusion. *Scientific Reports*, 9:1–17, 10 2019. doi: 10.1038/s41598-019-50168-2.

- P. Whiteley. Forecasting seats from votes in british general elections. *The British Journal of Politics International Relations*, 7:165 – 173, 05 2005. doi: 10.1111/j.1467-856X.2005.00179.x.
- R. Wilson. A Bidding Model of Perfect Competition. *Review of Economic Studies*, 44(3):511–518, 1977. URL <https://ideas.repec.org/a/oup/restud/v44y1977i3p511-518..html>.
- R. Wilson. Strategic analysis of auctions. In R. Aumann and S. Hart, editors, *Handbook of Game Theory with Economic Applications*, volume 1, chapter 08, pages 227–279. Elsevier, 1 edition, 1992. URL <https://EconPapers.repec.org/RePEc:eee:gamchp:1-08>.
- R. B. Wilson. Competitive bidding with disparate information. *Management Science*, 15(7):446–448, 1969. ISSN 00251909, 15265501. URL <http://www.jstor.org/stable/2628640>.
- J. Wolfers and E. Zitzewitz. Prediction markets. *Journal of Economic Perspectives*, 18:107–126, 02 2004.

Appendix A

Appendix To Chapter 3

A.0.1 Derivation of Monopoly Dealer Payoff Function

This can be solved in 2 ways: from first principles and from the general payoff integral. Here is the first principles derivation: Let M be the total number of clients and let there be an equal number of buyers and sellers. Let δ be the bid ask spread chosen by the monopoly dealer and let each client and the dealer receive a signal as to the true value of the asset which is normally distributed with mean V , the true value, and standard deviation σ_M for the dealer and σ_T for the clients.

$$E[\Pi] = E[\text{profits from buys}] + E[\text{profits from sells}]$$

$$E[\Pi] = \frac{M}{2} E[(\frac{\delta}{2} + \varepsilon)(1 - \Phi[\frac{2\varepsilon + \delta}{2\sigma_T}]) + (\frac{\delta}{2} - \varepsilon)(\Phi[\frac{2\varepsilon - \delta}{2\sigma_T}])]$$

$$E[\Pi] = \frac{M}{2} (\frac{\delta}{2} E[1 - \Phi[\frac{2\varepsilon + \delta}{2\sigma_T}]] + \frac{\delta}{2} E[\Phi[\frac{2\varepsilon - \delta}{2\sigma_T}]] - \frac{M}{2} \text{cov}(\varepsilon, \Phi[\frac{2\varepsilon + \delta}{2\sigma_T}]) - \frac{M}{2} \text{cov}(\varepsilon, \Phi[\frac{2\varepsilon - \delta}{2\sigma_T}]))$$

$$E[\Pi] = \frac{M}{2} (\frac{\delta}{2} [1 - \Phi[\frac{\frac{\delta}{2}}{2\sigma_T \sqrt{1 + \frac{\sigma_M^2}{\sigma_T^2}}}]]) + \frac{\delta}{2} \Phi[\frac{-\frac{\delta}{2}}{2\sigma_T \sqrt{1 + \frac{\sigma_M^2}{\sigma_T^2}}}] - \frac{M}{2} \text{cov}(\varepsilon, \Phi[\frac{2\varepsilon + \delta}{2\sigma_T}]) - \frac{M}{2} \text{cov}(\varepsilon, \Phi[\frac{2\varepsilon - \delta}{2\sigma_T}]))$$

$$E[\Pi] = \frac{M}{2} (\delta [1 - \Phi[\frac{\frac{\delta}{2}}{2\sqrt{\sigma_T^2 + \sigma_M^2}}]]) - M \text{cov}(\varepsilon, \Phi[\frac{2\varepsilon + \delta}{2\sigma_T}])$$

$$\text{cov}(\varepsilon, \Phi[\frac{2\varepsilon + \delta}{2\sigma_T}]) = E_\varepsilon[\Phi'[\frac{2\varepsilon + \delta}{2\sigma_T}]] \text{cov}(\varepsilon, \varepsilon) = \frac{\sigma_M^2}{\sigma_T} E_\varepsilon[\phi[\frac{2\varepsilon + \delta}{2\sigma_T}]]$$

by Steins' Lemma, where $\varepsilon \sim N(0, \sigma_M)$

$$= \frac{\sigma_M^2}{\sigma_T} \int_{-\infty}^{\infty} \phi[\frac{2t + \delta}{2\sigma_T}] \cdot \phi(t) dt = \frac{\sigma_M}{\sigma_T} \frac{1}{2\pi} \int_{-\infty}^{\infty} e^{-\frac{1}{2}(\frac{2t + \delta}{2\sigma_T})^2} e^{-\frac{1}{2}(\frac{t}{\sigma_M})^2} dt$$

let $b = \frac{\delta}{2\sigma_T}$

$$= \frac{\sigma_M}{\sigma_T} \frac{1}{2\pi} \int_{-\infty}^{\infty} e^{-\frac{1}{2}(\frac{t^2}{\sigma_T^2} + \frac{2b}{\sigma_T}t + b^2 + \frac{t^2}{\sigma_M^2})} dt = \frac{\sigma_M}{\sigma_T} \frac{1}{2\pi} e^{-\frac{b^2}{2}} \int_{-\infty}^{\infty} e^{-\frac{1}{2}(t^2(\frac{1}{\sigma_T^2} + \frac{1}{\sigma_M^2}) + \frac{2b}{\sigma_T}t)} dt$$

Let $s = (\frac{1}{\sigma_T^2} + \frac{1}{\sigma_M^2})^{0.5}t$, so $dt = (\frac{1}{\sigma_T^2} + \frac{1}{\sigma_M^2})^{-0.5}ds$

$$= \frac{\sigma_M}{\sigma_T} \frac{1}{2\pi} e^{-\frac{b^2}{2}} \int_{-\infty}^{\infty} e^{-\frac{1}{2}(s^2 + \frac{2b}{\sigma_T}(\frac{1}{\sigma_T^2} + \frac{1}{\sigma_M^2})^{-0.5}s)} (\frac{1}{\sigma_T^2} + \frac{1}{\sigma_M^2})^{-0.5} ds$$

$$= \frac{\sigma_M}{\sigma_T} \frac{1}{2\pi} e^{-\frac{b^2}{2}} (\frac{1}{\sigma_T^2} + \frac{1}{\sigma_M^2})^{-0.5} \int_{-\infty}^{\infty} e^{-\frac{1}{2}(s + \frac{b}{\sigma_T}(\frac{1}{\sigma_T^2} + \frac{1}{\sigma_M^2})^{-0.5})^2 - \frac{b^2}{\sigma_T^2}(\frac{1}{\sigma_T^2} + \frac{1}{\sigma_M^2})^{-1}} ds$$

Let $t = s + \frac{b}{\sigma_T}(\frac{1}{\sigma_T^2} + \frac{1}{\sigma_M^2})^{-0.5}$

$$\frac{\sigma_M}{\sigma_T} \frac{1}{2\pi} e^{-\frac{b^2}{2} + \frac{1}{2} \frac{b^2}{\sigma_T^2} (\frac{1}{\sigma_T^2} + \frac{1}{\sigma_M^2})^{-1}} (\frac{1}{\sigma_T^2} + \frac{1}{\sigma_M^2})^{-0.5} \int_{-\infty}^{\infty} e^{-\frac{1}{2}(s + \frac{b}{\sigma_T}(\frac{1}{\sigma_T^2} + \frac{1}{\sigma_M^2})^{-0.5})^2} ds$$

$$= \frac{\sigma_M}{\sigma_T} \frac{1}{\sqrt{2\pi}} e^{-\frac{b^2}{2} + \frac{1}{2} \frac{b^2}{\sigma_T^2} (\frac{1}{\sigma_T^2} + \frac{1}{\sigma_M^2})^{-1}} (\frac{1}{\sigma_T^2} + \frac{1}{\sigma_M^2})^{-0.5}$$

$$= \frac{\sigma_M}{\sigma_T} \frac{1}{\sqrt{2\pi}} e^{-\frac{\delta^2}{8\sigma_T^2} (1 - \frac{1}{\sigma_T^2} (\frac{\sigma_T^2 \sigma_M^2}{\sigma_T^2 + \sigma_M^2})} (\frac{1}{\sigma_T^2} + \frac{1}{\sigma_M^2})^{-0.5} = \frac{\sigma_M}{\sigma_T} \frac{1}{\sqrt{2\pi}} \frac{\sigma_T \sigma_M}{\sqrt{\sigma_M^2 + \sigma_T^2}} e^{-\frac{\delta^2}{8(\sigma_T^2 + \sigma_M^2)}}$$

$$= \frac{1}{\sqrt{2\pi}} \frac{\sigma_M^2}{\sqrt{\sigma_M^2 + \sigma_T^2}} e^{-\frac{\delta^2}{8(\sigma_T^2 + \sigma_M^2)}}$$

$$\therefore E[\Pi] = \frac{M}{2} \delta (1 - \Phi[\frac{\delta}{2\sqrt{\sigma_T^2 + \sigma_M^2}}]) - M \frac{\sigma_M}{\sigma_T} \frac{1}{\sqrt{2\pi}} \frac{\sigma_T \sigma_M}{\sqrt{\sigma_M^2 + \sigma_T^2}} e^{-\frac{\delta^2}{8(\sigma_T^2 + \sigma_M^2)}}$$

This is expressed in terms of the bid ask spread, so as a function of the semi bid ask spread:

$$E[\Pi(\delta, \sigma_T, \sigma_M)] = M \delta (\Phi[\frac{-\delta}{\sqrt{\sigma_T^2 + \sigma_M^2}}]) - \frac{M}{\sqrt{2\pi}} \frac{\sigma_M^2}{\sqrt{\sigma_M^2 + \sigma_T^2}} e^{-\frac{\delta^2}{2(\sigma_T^2 + \sigma_M^2)}} \quad \square$$

Appendix B

Appendix To Chapter 4

B.0.1 Formal Specification of Static Bayesian Games

In a Bayesian game, we specify strategy spaces, type spaces, payoff functions and beliefs for every player. As in perfect information games, a strategy is a complete plan of action that covers every possibility that might occur for every player type that might occur. A type space is the set of all possible types for a player

Like a game of complete information, a game of incomplete information has (i) a set of players and (ii) their action spaces. These are complemented with preferences and information components: (iii) a probability distribution over players' types which determine their preferences, (iv) each player knows his own type but not the other players' types; (v) the probability distribution over types is common knowledge, which is itself common knowledge; (vi) Payoffs associated with each action space and type space.

Before the game is played, Nature chooses the different player types. Each type t_i can represent information about player i 's own payoffs, or more generally, other game attributes in particular, network structure. Thus there is a type space T_i for each player $i \in N$, representing the range from which Nature chooses i 's type. We introduce a commonly known prior probability distribution $p(\cdot)$ on $\prod_{i=1}^N T_i$ to describe how Nature chooses a type profile $(t_i)_{i=1}^N$

The normal form representation of a n -player static Bayesian game is:

$$G = \langle N, (A_i)_{i=1}^n, (T_i)_{i=1}^n, ((u_i(\cdot; t_i))_{t_i \in T_i}, (p_i)_{i=1}^n) \rangle$$

where $N = \{1, 2, \dots, n\}$ is the set of players A_i is player i 's action set $T_i = \{t_i^1, t_i^2, \dots, t_i^{k_i}\}$

is player i 's type space, and $u_i : A \times T_i \rightarrow \mathfrak{R}$ is player i 's type dependent utility or payoff function, where $A := A_1 \times A_2 \times \dots \times A_n = \prod_{i=1}^n A_i$

Timing of the Bayesian game:

1. Nature chooses a type profile $(t_i)_{i=1}^n \in \prod_{i=1}^n T_i$
2. Each player $i \in N$ learns his own type $t_i \in T_i$ which is his private information, and uses his prior p_i to form beliefs $p_i(t_{-i}|t_i)$ over the other player types
3. Players move simultaneously and choose actions $a_i \in A_i$ and payoffs $u_i(a, t)$ for $i \in N$ are realized

A pure strategy for player i is a function $f_i : T_i \rightarrow A_i$ that specifies a pure action $f_i(t_i)$ that i will choose when his type is t_i

Types may also be derived from continuous distributions, with the random type space T_i with a CDF $F_i(t_i)$ and density $f_i(t_i)$

In the static Bayesian game, $G = \langle N, (A_i)_{i=1}^n, (T_i)_{i=1}^n, ((u_i(\cdot; t_i))_{t_i \in T_i}, (p_i)_{i=1}^n) \rangle$ a pure strategy profile $s^* = (s_1^*(\cdot), s_2^*(\cdot), \dots, s_n^*(\cdot))$ is a pure strategy Bayesian Nash Equilibrium if for every player $i \in N$, and for each realization $t_i \in T_i$ of player i 's type, the action $a_i = s_i^*(t_i)$ is a best response because it solves:

$$\max_{a_i \in A_i} E_{t_{-i}}[u_i(a_i, s_{-i}^*(t_{-i}); t_i) | t_i]$$

We take player i 's conditional expectations over the random realizations of other player types t_{-i} , given that player i knows his own type t_i

All finite Bayesian games have a Nash equilibrium.

B.0.2 Convergence with an Agent Based Model (ABM)

As a second check, we recreated the OTC auction game in an ABM, where each of the clients quotes and trades with a set of dealers, and the dealers attempt to maximize their individual payoffs. It uses a modified version of the relaxation algorithm where dealers do not know their payoff function but can measure the effect on their own payoff by changing spreads. The results are in line with Relaxation algorithm solutions to the simultaneous dealer payoff functions.

An ABM was created and each dealer tries a higher or lower spread and compares the average payoff. If the high or lower spreads produce a higher average, then that




FIGURE B.1: convergence of dealers' equilibrium spreads in ABM

spread is adopted instead. All agents perform this operation in turn, similar to the methodology of the relaxation algorithm. Figure B.1 shows the convergence of the bid ask spreads in an Agent Based Model producing a mean spread 0.7525 with sd of 0.0071 vs theoretical Nash equilibrium level of 0.75. Dealers find the equilibrium after approximately 50 rounds of trading auctions with no knowledge of the payoff functions