

ESSAYS IN EMPIRICAL INDUSTRIAL ORGANIZATION

Mateusz Myśliwski

University College London (UCL)

Thesis submitted in partial fulfilment of the degree of

Doctor of Philosophy (PhD) in Economics

London, May 2019

Declaration

I, Mateusz Myśliwski, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

Mateusz Myśliwski

London, May 2019

Abstract

This thesis comprises three essays addressing questions related to pricing, competition, and consumer welfare in imperfectly competitive industries characterised by the existence of demand and supply side frictions.

Chapter 1 studies a dynamic pricing game amongst differentiated multiproduct oligopolists who have incentives to temporarily lower prices to attract new consumers who are more likely to purchase the same product again at a higher price due to inertia. The novel feature of the model, which allows to explain persistence in the observed patterns of retail prices, is the possibility of costly price adjustment. The magnitude of adjustment costs is estimated using scanner data on purchases of a category of dairy products by UK households. The main results suggest that adjustment costs can be substantial for manufacturers, but they are passed through to the consumers only on a very limited scale.

Chapter 2 explores an alternative explanation for price dispersion by introducing a model with search frictions and private information about marginal costs. Consumers decide how many sellers to visit who in equilibrium set prices in a fashion similar to bidding in a reverse first-price auction with an unknown number of competitors. The chapter shows how the distributions of search costs and firm heterogeneity can be nonparametrically identified and analyses convergence rates of the proposed estimators.

Chapter 3 uses an extension of the search model proposed earlier to study the value of information provided by mortgage brokers in the UK. Prospective borrowers can either directly search and apply for mortgages with different lenders or use a broker who finds the best rate on their behalf. The main finding suggests that, on average, the existence of brokers substantially fosters competition between lenders, leads to lower monthly payments in equilibrium and reduces the deadweight loss arising as a result of costly search.

Impact Statement

The research presented in the chapters of the thesis can broadly impact the academic and non-academic discussion on competition, regulation, market power, and consumer welfare in imperfectly competitive markets. The main results highlight new channels through which oligopolists' market power can be reduced, even if consumers themselves exhibit inertia (chapter 1) or information acquisition is costly (chapter 3).

The empirical findings in chapter 1 are of particular interest for regulators who have been raising concerns about the lack of transparency in the payments between manufacturers and retailers in vertical chains. The framework presented there allows to estimate the magnitude of promotional fees borne by manufacturers to influence pricing decisions by the downstream firm. These payments are typically hidden in the contracts and their existence has been a subject of a widespread debate in the media. The chapter shows that removing such payments would indeed lead to a redistribution of profits within the channel, but would not lead to a big improvement in terms of consumer surplus. These findings can be used in the design of policies regulating contracts between manufacturers and retailers.

The conclusions in chapter 3 contribute to the important discussion on whether markets benefit from intermediation. I show that in industries where direct contact between buyers and sellers (here borrowers and lenders) can be difficult because of search frictions, the existence of brokers positively affects competition and leads to lower prices. However, an important finding from this chapter, which should be a pivotal argument in the public debate, is that the effects of intermediation are heterogeneous across different types of consumers and products they choose. Therefore, no regulation should be uniformly affecting every party in the same way. In the context of the UK mortgage market, certain types of borrowers could be encouraged to obtain advice from brokers more often than others.

The models developed in the chapters of the thesis have novel features and can be used to study a wider array of industries. The methodology proposed in chapter 1 can be readily extended to study dynamic pricing problems with an arbitrary number of firms, products and consumer types. While

discussing the search model in chapter 2, I put particular emphasis on explaining the estimation strategy and potential extensions, hoping that other researchers can benefit from the methodological innovations introduced in these chapters.

Parts of the research in this thesis have already been disseminated to a wider audience through presentations at international conferences and invited seminars. Finally, since the three chapters can be treated as self-contained papers, my important goal for the future is to publish them in leading academic journals.

Acknowledgements

"But anything to be written has to be, time and again, begun from the start, and time and again attempted anew, until one day it succeeds at least approximately, if never quite satisfactorily. No matter how unpromising it is and no matter how terrible and hopeless, if we have a subject which time and again, and yet time and again, grips us with the utmost persistence and no longer leaves us alone, it should time and again be attempted. In the knowledge that nothing at all is certain and that nothing at all is perfect, we should, even with the greatest uncertainty and with the greatest doubts, begin and continue whatever we have determined to do. If we give up each time even before we have started, we eventually find ourselves in desperation, and finally and ultimately we no longer get out of that desperation and are lost."

"Just as we wake every day and have to begin and continue what we have determined to do, that is to continue existing, quite simply because we have to continue existing, so we must begin and continue such an enterprise (...)."

Thomas Bernhard "Yes", translated from German by Ewald Osers

My PhD studies and this thesis would not have been possible without the funding I received from multiple sources. I would like to thank the ESRC and UCL for supporting me with the Andrew Szmidla Postgraduate Scholarship in my first year.

Moreover, I would like to thank the Bank of England for allowing me to work on the third chapter of the thesis during my internship by providing access to their facilities and, most importantly, the data. Any views expressed in that chapter are solely those of the author and so cannot be taken to represent those of the Bank of England or to state Bank of England policy. This chapter does not necessarily represent the views of the Bank of England or members of the Monetary Policy Committee or Financial Policy Committee. All errors are my own.

The empirical analysis in chapter 1 used data supplied by Kantar Worldpanel obtained with financial support from the ERC under ERC-2009-AdG grant agreement number 249529. The use of Kantar Worldpanel data in this work does not imply the endorsement of Kantar Worldpanel in relation to the interpretation or analysis of the data.

On a more personal note, I am grateful to Lars Nesheim for his guidance and supervision throughout my PhD. I have also benefitted from discussions with many faculty members of the Department of Economics at UCL, in particular Áureo de Paula and Nikita Roketskiy.

I am especially indebted to my co-authors, May Rostom, Fabio Sanches, Daniel Silva, and Tang Srisuma, who not only generously allowed me to use our joint work in my thesis, but also were always ready to discuss and improve my ideas and push the papers in the right directions.

The emotionally exhausting process that led to the completion of this thesis was made significantly more bearable thanks to my friends from the PhD. There are too many of you to mention separately, but everyone who ever spent their time in Drayton G06 should feel acknowledged.

Last but not least, none of it would have happened if it were not for my family in Poland and my partner and the love of my life Kasia, who has stuck with me through thick and thin. This thesis is dedicated to you.

Contents

Preface	1
1 Implications of Consumer Loyalty for Price Dynamics when Price Adjustment is Costly	5
1.1 Introduction	5
1.2 Related literature	8
1.3 Data and industry background	12
1.4 Descriptive evidence	17
1.4.1 Consumer loyalty	17
1.4.2 Price rigidities	19
1.4.3 Implications of consumer loyalty and price adjustment costs for price dynamics	20
1.5 Model	23
1.5.1 Preliminaries	24
1.5.2 Firms	24
1.5.3 Consumers	26
1.5.4 Equilibrium	29
1.6 Identification and estimation	29
1.6.1 Identification	29
1.6.2 Estimation	31
1.6.3 Dynamic parameters	35
1.7 Model fit and counterfactuals	39
1.7.1 Model fit	40
1.7.2 Counterfactuals	41
1.8 Summary and conclusions	45
2 Identification and Estimation of a Search Model: A Procurement Auction Approach	47
2.1 Introduction	47

2.2	Model	51
2.2.1	Supply Side	51
2.2.2	Demand Side	52
2.2.3	Equilibrium	53
2.3	Identification	54
2.3.1	Demand Side	54
2.3.2	Supply Side	55
2.3.3	Constructive Identification	56
2.4	Main Results	57
2.4.1	Nonparametric Restrictions on the Data	57
2.4.2	Large Sample Properties	58
2.5	Possible Extensions	62
2.5.1	Relaxing Assumption I	62
2.5.2	Asymmetric Search Probabilities	63
2.6	Simulation	64
2.7	Concluding Remarks	66

3 Value of Information and the Impact of Mortgage Intermediaries on Lender Competition and Households' Financial Positions 68

3.1	Introduction	68
3.2	The UK mortgage market and data	72
3.2.1	Institutional overview	73
3.2.2	Data	74
3.2.3	Mortgage cost	75
3.2.4	Reduced form findings	75
3.3	Model	79
3.3.1	Borrowers	80
3.3.2	Lenders and equilibrium	81
3.4	Identification and estimation	83
3.4.1	Differences from Salz (2017)	84
3.4.2	Role of exclusion restrictions	85
3.4.3	Estimation steps	86
3.5	Results	88
3.5.1	Borrowers' search costs	90

3.5.2	Lenders' costs and margins	92
3.6	Counterfactuals	95
3.6.1	Value of information provided by brokers	95
3.6.2	Market centralisation	99
3.6.3	Summary of findings	100
3.7	Conclusion	101
	Concluding remarks	103
	Appendices	
	Appendix A Supplementary material for chapter 1	104
A.1	Alternative demand specification	104
A.2	Main identification result	105
A.3	Identification of promotional costs	109
A.4	Estimation of the discount factor	114
A.5	Model solution	116
A.6	Additional tables	118
A.7	Additional results	121
	Appendix B Supplementary material for chapter 2	126
B.1	Proofs	126
	Appendix C Supplementary material for chapter 3	128
C.1	Data and reduced form results	128
C.1.1	Data and summary statistics	128
C.1.2	Estimation sample	130
C.1.3	Probability of using a broker	130
C.1.4	Robustness checks	133
C.2	Additional figures	136
	References	138

List of Figures

1.1	Prices of 500g spreadable butters and margarines in Tesco stores.	16
1.2	Price promotions and market shares of Country Life (500g spreadable) in Tesco.	16
2.1	Monte Carlo performance of kernel density estimators for the distribution of prices.	65
2.2	Monte Carlo comparison of two infeasible estimators.	65
2.3	Monte Carlo comparison of two infeasible estimators.	66
2.4	Monte Carlo performance of the feasible density estimator.	66
3.1	Equilibrium sorting.	82
3.2	Estimated search cost CDFs.	91
3.3	Marginal distributions of marginal costs.	94
3.4	Distribution of price-cost margins.	94
3.5	Distributions of prices and price-cost margins in the counterfactual scenarios.	101
A.6.1	Histograms with the number of products on sale by firm.	120
A.7.1	Actions played by firms: model vs. data.	124
A.7.2	Market shares by brand: model vs. data.	125
C.1.1	Distributions of predicted probabilities of using a broker.	132
C.2.1	Distributions of price-cost margins.	136
C.2.2	Pairwise comparisons of estimated search cost CDFs.	137

List of Tables

1.1	Expenditure shares of main supermarket chains in the butter and margarine category.	13
1.2	Manufacturers and brands of top 6 branded products.	15
1.3	Consumer switching patterns for purchases made in two subsequent weeks. . .	18
1.4	Consumer switching patterns	18
1.5	Frequency of price changes.	19
1.6	Durations of promotions in Tesco stores.	20
1.7	Price regressions	22
1.8	Demand estimates	33
1.9	Price adjustment costs.	36
1.10	Estimated discount factors.	38
1.11	Magnitude of adjustment costs.	39
1.12	Measures of model fit.	40
1.13	Implications of consumer loyalty with and without price adjustment costs. . . .	42
1.14	Counterfactual results with $SC = 0$	44
1.15	Decomposition of main counterfactual results.	45
3.1	Price dispersion by sales channel.	77
3.2	Price benefits of using a broker.	77
3.3	Impact of using a broker on price plus broker fees.	78
3.4	Covariate selection.	89
3.5	Summary statistics for estimated search cost distributions.	90
3.6	Summary statistics for estimated marginal costs and margins.	93
3.7	Value of information: breakdown by borrower types.	97
3.8	Value of information: breakdown by loan characteristics.	98
3.9	Price changes in a centralised market.	100
A.6.1	Annual expenditure shares in the 500g spreadable segment.	118

A.6.2	Price levels.	119
A.7.1	Demand estimates (alternative definition of loyalty).	121
A.7.2	Demand estimates (heterogenous γ).	122
A.7.3	Multinomial logit CCP estimates.	123
C.1.1	Summary statistics	128
C.1.2	Interest Rates by LTV and Rate Duration	129
C.1.3	Share by Household type and LTV	129
C.1.4	Share by Bank and Product Type	130
C.1.5	Raw and Final Sample	130
C.1.6	Probability of using a broker	131
C.1.7	Price benefits of using a broker: brokers who do not charge the borrowers.	133
C.1.8	Price benefits of using a broker: brokers who are not paid by lenders.	133
C.1.9	Value of information with adjusted costs: breakdown by borrower types.	134
C.1.10	Value of information with adjusted costs: breakdown by loan characteristics.	135

Preface

The last three decades have seen a steady increase in the number of empirical studies of imperfectly competitive industries, which use the structure of game-theoretic economic models and modern econometric tools to analyse market data and answer various questions related to competition, market structure and industry dynamics (Akerberg et al., 2007). However, the business environment in many industries changes so rapidly that even with an increased availability and improved quality of data, IO researchers constantly need to broaden the set of tools and models they use, in order to be able to answer questions which are up-to-date and relevant for academics and policy-makers. Therefore, the majority of research presented in this thesis was motivated by the need to fill some methodological gaps in the empirical IO literature and then applying the proposed methods to answer questions using detailed data from two important industries: retailing and banking. In particular, I was interested in formulating models that: (i) would be able to detect and measure the magnitude of different types of supply and demand frictions including price adjustment costs faced by firms, consumer inertia, and search costs; (ii) could be used to counterfactually analyse how such frictions affect pricing dynamics, competition and consumer welfare.

Chapter 1 is based on my joint paper with Fabio Sanches, Daniel Silva and Tang Srisuma and was inspired by my earlier work on dynamic discrete games. While analysing data on dynamics of retail prices of consumer packaged goods, it became clear that for most products, researchers observe only a discrete set of prices which also exhibit a significant degree of persistence over time.

To explain such patterns, the chapter introduces a new dynamic game-theoretic model in which multiproduct firms simultaneously choose prices for all differentiated products they offer, which can be structurally estimated using panel data on consumers' choices and observed prices. While the players involved in the game are forward-looking, consumers are myopic and exhibit brand loyalty. This generates a dynamic pricing incentive for the firms, which face the well-documented trade-off between building up a consumer base by charging low prices today and exploiting the fact that some consumers are already locked in with the product and therefore they are willing to

pay a higher price now. The source of frictions on the supply side are the costs of adjusting prices. In other words, even when firms want to change the prices, they have to factor in the administrative expenses they have to incur. Basing on the aforementioned observation that prices in our data take only a limited number of values, we depart from the traditional treatment of prices as continuous decision variables, in favour of a discrete choice approach. The additional merit of following this path is that the model turns out to be a particular instance of the dynamic oligopoly framework of Ericson and Pakes (1995) and therefore we can use the computationally feasible two-step estimation techniques developed for dynamic entry games. Moreover, after introducing firm-level private information, we are guaranteed that an equilibrium in pure strategies exists, which has been a problematic issue in papers studying optimal pricing with consumer inertia, such as Dubé, Hitsch, and Rossi (2009) and Pavlidis and Ellickson (2017).

The features of the model are of particular importance for the empirical application to the UK butter and margarine industry. This is a typical example of an oligopoly with three dominant firms who sell multiple products under different brand names. The only dimension of price competition in this market is through temporary price promotions (sales), and more specifically switching between regular and sale prices (Hosken and Reiffen, 2004). Industry reports reveal that 70% of supermarket suppliers make payments toward marketing costs or price promotions. Moreover, since promotional fees are usually not regulated by long-run contracts and instead negotiated on a running basis, they have attracted interest from the UK competition authorities. Our research aims to confront some of the anecdotal claims mentioned above with data. Firstly we seek to quantify the magnitude of promotional costs to see whether they constitute an important part of manufacturers' revenues. We then analyse how imposing regulation on the payments from manufacturers to retailers would affect consumer welfare and firms' profits.

The contribution of chapter 2, based on another joint paper with Fabio Sanches, Daniel Silva and Tang Srisuma, is mostly methodological. The chapter introduces a new model of nonsequential search with a continuum of consumers and a finite number of firms. Both consumers and firms are heterogeneous. Consumers differ in search costs. Firms have private marginal costs of production. This assumption is different from virtually all structural models of search, starting with Hong and Shum (2006), which, following Burdett and Judd (1983) assumed that firms are ex ante identical and play mixed strategy in equilibrium. We show that an equilibrium price dispersion can arise in this model both as a result of search and ex ante heterogeneity amongst firms. However, heterogeneity in cost is not sufficient and some degree of search is still necessary for firms not to charge monopoly price to all consumers in equilibrium. We provide conditions to identify the features

of the model using prices and market shares. Our identification strategy is constructive and relies on results from the literature on nonparametric auction estimation. We derive the uniform rate of convergence of our estimator and show its properties in a series of Monte Carlo simulations.

The final chapter of the thesis, joint with May Rostom, is an empirical study of the impact of search frictions and brokerage services on mortgage choices and competition in the UK credit industry. We extend the structural model of search with lender and borrower heterogeneity proposed in chapter 2 to estimate the value of information provided by mortgage intermediaries (brokers) in the UK. Using administrative data on loans originated in 2016 and 2017, we document the existence of a substantial degree of unexplained price dispersion and observe that while mortgages obtained through brokers are on average cheaper, borrowers who use intermediaries end up paying more once commissions are factored in. This fact underpins the assumption that brokers are used by borrowers with higher search costs, which helps nonparametrically identify and estimate the distributions of search cost and banks' cost of providing the loan. Our results show that search costs vary by demographic groups and brokers' presence exerts a negative pressure on lenders' market power. To estimate how intermediation affects consumer surplus, we consider a counterfactual where broker advice is not available, finding that brokers' presence reduces average monthly costs by 33.7% and welfare losses caused by search frictions by 16.5%, though the results differ by borrower and loan characteristics. Remarkably, average value of information provided by brokers is positive only for borrowers who take up mortgages with 2-year fixed term deals. In a second counterfactual we look at the effects of a hypothetical market centralization, finding that such a regulation would halve lenders' markups and lower monthly costs of an average mortgage by 6.4%.

Last part of the thesis includes concluding remarks and outlines directions for future research, both methodological and empirical. Additional materials for each of the chapters, including omitted derivations, proofs, and additional results are included in the appendices.

Chapter 1

Implications of Consumer Loyalty for Price Dynamics when Price Adjustment is Costly

1.1 Introduction

A general consensus in the marketing and industrial organisation literature is that the existence of consumer switching costs, due to brand loyalty (habit or other types of inertia), creates two countervailing effects for firms' pricing decisions: *investing* and *harvesting*. See, for example, Beggs and Klemperer (1992), Dubé, Hitsch, and Rossi (2009), Arie and Grieco (2014). The investment motive acts as an incentive for firms to temporarily lower their prices in order to build up a larger base of loyal consumers. Subsequently, after acquiring a number of loyal consumers, firms may increase their prices and harvest the investments made in the previous periods. Existing theoretical models of competition do not deliver unequivocal predictions about the effect of consumer loyalty on equilibrium prices, concluding that they can be either higher (if the harvesting motive prevails) or lower (if the opposite is true) (Farrell and Klemperer, 2007). To supplement theoretical findings, recent empirical studies by Dubé et al. (2009) and Pavlidis and Ellickson (2017) have shown that the investing motive tends to generate a negative pressure on prices.

All models that study the effects of consumer loyalty on price dynamics assume that firms can adjust their prices freely. This stands in contrast to a substantial body of empirical research in economics that highlights the importance of price adjustment costs from the supply side in a variety of settings. For example, see Slade (1998), Aguirregabiria (1999), Dutta, Bergen, Levy, and Venable (1999), Levy, Bergen, Dutta, and Venable (1997), Zbaracki, Ritson, Levy, Dutta, and Bergen

(2004) and Ellison, Snyder, and Zhang (2015). If price adjustments are in fact costly, intuitively, when firms reduce prices to invest in consumer loyalty they not only have to deal with temporary profit losses but also with the cost of the price change itself. As a consequence, price adjustment costs may alter firms incentives to invest in new consumers, with potential repercussions on the correlation between loyalty and equilibrium prices.

In this chapter we present an empirical model of dynamic oligopoly pricing that explicitly models both demand and supply side frictions. We use it to study the UK butter and margarine industry using a rich scanner dataset. We are interested in knowing whether price adjustment costs are present for the dairy producers who supply their products to supermarkets. If the answer is yes, we would want to investigate: (i) how they can affect prices through promotions, and (ii) their implications on firm profits and consumer welfare in the presence of consumer loyalty.

The focus on adjustment costs is of a significant interest in the context of our application. One interpretation of the adjustment costs in our model is the promotional fees that suppliers pay to the supermarkets; for example, for the purpose of prominently featuring their products on fliers or designated store shelves. These payments are key components of the so-called *supplier rebates* that, in the UK, are suspected to be substantial but hard to observe or quantify, even for financial accounting purposes.¹ Therefore counterfactual exercises that remove the adjustment costs can be used to make predictions following a ban on promotional fees.

The supply side of our model is based on a class of dynamic discrete games that has seen increasing number of applications in IO and other fields (e.g. Aguirregabiria and Mira (2007) or Pesendorfer and Schmidt-Dengler (2008)). The dairy industry is an example of an oligopoly with three dominant firms: Arla, Dairy Crest and Unilever, who sell multiple products under different brand names. Their main sales channels are national retail chains. Therefore price competition can occur through temporary price promotions (sales), and more specifically switching between regular and sale prices (Hosken and Reiffen, 2004). We therefore model price as a discrete variable. On the demand side consumers maximise their household utilities with their purchase options. We assume that consumers are myopic but exhibit some degree of brand loyalty. Our model is a particular instance of the dynamic oligopoly framework of Ericson and Pakes (1995) that is known to be computationally feasible and possess an equilibrium in pure strategies.

¹Supplier rebates for big supermarkets in the UK received some attention recently for their lack of transparency in firms' balance sheets. Unlike in the US, for example, UK retailers do not publish how much money they receive from commercial income. A BBC article published in October 2014 says that, according to Fitch, the declared income on supplier rebates from a number of big American supermarkets "are the equivalent to 8% of the cost of goods sold for the retailers, equal to virtually all their profit", and a chartered accountant who specialises in working with UK supermarket balance sheets "conservatively estimates supplier contributions to be worth around £5bn a year to the top four supermarkets". More details and discussions on supplier rebates can also be found at: <https://www.bbc.com/news/business-29629742> and <https://www.economist.com/business/2015/06/18/buying-up-the-shelves>.

Adjustment costs are generally not identified nonparametrically but can be identified under a normalisation (Aguirregabiria and Suzuki (2014), Komarova, Sanches, Silva Jr., and Srisuma (forthcoming)). Our normalisation choice is consistent with the motivation that suppliers bear the fees to sponsor a price promotion.² More specifically our application assumes a cost is incurred to the supplier when a product goes on promotion but there is no cost when it returns to its regular price. We otherwise allow the adjustment costs to be fully heterogenous across brands and supermarkets. We find that price adjustment costs are substantial in magnitude and constitute between 24% and 34% of firms' variable profits. In absolute terms, these estimates are very similar across players and given that the firms we considered are the market leaders, this result may indicate that price adjustment costs constitute a much bigger fraction of the profits of smaller companies and local dairies, effectively restricting the scope of their promotional activities. This is consistent with what we observe in the data on smaller producers, who put their products on promotion much less frequently. Our results also complements findings from the marketing literature that market shares are positively correlated with the frequency of temporary price cuts.³

Our counterfactual studies compare equilibrium outcomes from models with and without adjustment costs. We first analyse the impacts of consumer switching costs on prices. We do this by comparing equilibrium prices across different consumer loyalty levels. Our model predicts that increases in consumer switching costs lead to increases in equilibrium prices. But this effect is significantly more pronounced in the model with price adjustment costs than without. In particular, our estimates show that a three fold increase in consumer switching costs may lead to a price increase that is up to 250% higher in the model with price adjustment costs vis-à-vis the price increase observed in the model without price adjustment costs. Therefore ignoring price adjustment costs can substantially underpredict the effects of consumer loyalty on prices.

Next, we consider the implications of price adjustment costs for firm profits and consumer surplus. One can interpret this investigation as a welfare analysis of a ban on promotional fees.⁴ We find that when price adjustment costs are excluded from the model profits increase substantially, between 50-70%, but consumer surplus goes up by only 0.4-3.3%. This happens because manufacturers pass only a small fraction of the cost reduction to the consumers. When we remove price adjustment costs from our model the frequency of promotions increases and the average duration

²See the last paragraph of the concluding section in Aguirregabiria (1999), as well as Blattberg and Briesch (2010) for a more detailed description of promotional mechanisms.

³For example, Agrawal (1996) noted that smaller brands should rather focus on advertising than price promotions. In the context of slotting fees, Bloom et al. (2000) established that the existence of payments from manufacturers to retailers might be hindering competition because these costs are higher for smaller brands in relative terms.

⁴This does not mean suppliers do not pay retailers for other costs. Operational costs, which include menu costs of printing new labels and organising shelves, can still enter firm profit functions.

of promotional spells decreases. In line with previous estimates (see Basker (2015) and Ellison et al. (2015)), our results indicate that investments in technologies that seek to reduce the costs of adjusting prices may generate considerable returns for firms. Alternatively, the results from this counterfactual can be interpreted as the welfare estimates of a ban on promotional fees.

Our estimation strategy combines different methodologies. We use household level scanner data to estimate a state-dependent logit demand model, and obtain a law of motion for aggregate market shares. The other components of the firm payoff functions are separated into the adjustment costs and everything else. We estimate the adjustment costs using the approach proposed in Komarova et al. (forthcoming), who show that switching costs in dynamic games – for example, entry costs in entry games, capacity adjustment costs in investment games, and promotional fees in the context of our application – can be identified in closed form. Furthermore, the estimates of adjustment costs are robust to different specification of profits and the discount factor. We also estimate the discount factor, which, as shown by Komarova et al. (forthcoming), is identified when period payoffs are linear in parameters, which is true in the case of the game analysed in this work. The estimated discount factors are between 0.92 and 0.99 for different suppliers, which lie within the range of values commonly assumed by other papers in this literature, suggesting that pricing decisions have an important intertemporal component.

The rest of the chapter is structured as follows: the next section provides a brief literature review, followed by several facts about the industry and description of the data. Section 1.4 presents some preliminary reduced-form evidence of consumer switching costs and price adjustment costs and their implications for pricing decisions. Section 1.5 introduces the theoretical model. Section 1.6 explains our identification strategy, steps of the estimation procedure, and shows our structural estimates. We then discuss the fit of our model and our main counterfactual results in section 1.7. Section 1.8 concludes the chapter. Appendix A contains derivations and additional details on the identification strategy, description of the algorithms used to solve the model and robustness checks to some of our assumptions.

1.2 Related literature

The model consists of three main building blocks: multiproduct nature of firms, consumer- and firm-side switching costs. To the best of our knowledge, there exist no theoretical or empirical papers analysing these three aspects jointly. We also make use of some predictions stemming from microeconomic models of sales and, in the empirical part of this work, of the recent developments in the literature on identification and estimation of dynamic games.

■ **Models of sales.** The existence of temporary sales is not a new phenomenon and pricing problems of this kind have been extensively studied, both in the economics and marketing literature.⁵ Interestingly, explanations for transient downward price movements substantially differ from one another, depending on the assumptions of the model and characteristics of industry or product of interest. Even though in our model we assume that sales arise because firms are aware of consumer inertia, it is still worth noting that there might be other explanations for the observed patterns of prices, valid in other contexts.

Early theoretical contributions typically considered sellers of a homogenous good who offered temporarily lower prices to discriminate between informed and uninformed (Varian, 1980) or high- and low-valuation consumers (Conlisk et al., 1984). Sobel (1984) extended the model in the latter paper to an oligopolistic setting to find that there are equilibria in which firms act as monopolists to their loyal consumers and cut their prices only to compete for the price-sensitive consumers. Some other equilibria characterised by long spells of regular prices might arise in the presence of punishment strategies.⁶ Although we cannot account for history-dependence in the Markovian framework and the author considers a durable good in a very stylised environment, some insights from this study should also pertain to more complex situations, such as the one described in our work.

Aguirregabiria (1999) provides an alternative explanation for cyclical pricing based on retailers' inventory dynamics and costs associated with placing orders and changing prices. This is one of the first papers to use a structural model to produce an estimate of menu costs and notice that from retailers' perspective, the costs of cutting prices are relatively low, since it is the manufacturers (wholesalers) who bear most of it. Therefore, the estimates we obtain should complement the findings of Aguirregabiria (1999), insofar as we analyse the pricing problem from the manufacturers' standpoint. Pesendorfer (2002) attempts to explain pricing patterns of two ketchup brands using a model of intertemporal pricing. The main source of price dynamics in his paper is demand accumulation due to consumer stockpiling and purchase acceleration in the low-price periods. In order to test the implications of his theoretical model, the author runs a series of reduced form regressions, in which he treats the pricing decision as a discrete choice problem, just like we do in our setting. Another paper exploring the stockpiling side of the story is Hendel and Nevo (2013), who, in addition to confirming some of the previous findings in the literature, are able to quantify

⁵It is important to note, that we consider price discounts as the only dimension of firms' promotional activity in that paper. There are other studies which use the label *promotions* to talk about advertising which increases the demand faced by the firm. See for example the pioneering study by Schmalensee (1976).

⁶See also Nava and Schiraldi (2014) for another collusion-based explanation for sales.

the welfare effects of sales in the market for soft drinks. They find the total effects to be positive, since the manufacturers are able to partially mimic the otherwise unattainable third degree price discrimination. The net impact of sales on consumer welfare, on the other hand, turns out to be ambiguous: price-sensitive households tend to benefit from the temporary price markdowns, whereas households that are non-storers experience a decrease in welfare, relative to a uniform pricing scenario.⁷

The body of theoretical literature trying to explain why sales occur is still growing. Some recent papers try to broaden the scope of potential explanations for sales by including sellers' rational inattention (Matějka, 2016) and consumers' loss aversion (Heidhues and Kőszegi, 2014). These papers, together with the ones discussed in the previous paragraph, highlight that there are multiple ways to explain why temporary sales occur, and, confronted with the data, one usually finds mixed evidence when it comes to evaluating welfare effects.

■ **Multiproduct oligopoly.** Even though multiproduct firms are a widespread phenomenon now, economic theory still provides mixed guidance regarding the consequences of existence of such firms for welfare and market power. One important fact about multiproduct oligopolies is that not all of the theoretical implications of traditional single-product models pertain to settings with multiple differentiated products (see e.g. McAfee (1995), Johnson and Myatt (2006)). Analysing market structure in the presence of multiproduct firms, Anderson and De Palma (2006) found that since the product lines are usually too narrow and profits higher than with single-product firms, multiproduct environments encourage entry. Multiproduct pricing has been recently analysed by Shelegia (2012), who found that in a world characterised by information frictions, in a static equilibrium, prices of substitutable goods offered by multiproduct firms are uncorrelated and firms play mixed strategies. Rhodes (2014) analysed multiproduct retailing and arrived at a striking result that with two multiproduct firms selling the same two goods and consumer search, firms randomly choose whether to hold a sale, on which product and by how much the price should be reduced. In the empirical literature, Nevo (2001a) found price-cost margins emerging from a multiproduct Nash-Bertrand pricing consistent with data from the breakfast cereal industry. These margins are found to be lower than ones that would result from collusive pricing, but higher than in the case of single-product firms. The reason for this is that multiproduct firms are able to internalise part of the competition effects. A recent paper that is perhaps most closely related to our study is Pavlidis and Ellickson (2017), who consider oligopolists' dynamic pricing strategies

⁷When the good is storable, Hendel, Lizzeri, and Roketskiy (2014) show that cyclical patterns of high and low prices might arise even when consumers are homogenous.

in the presence of consumer switching costs and *umbrella branding* in the yogurt industry, that is selling multiple differentiated products under the same brand name. The key differences between their paper and our approach are that they do not consider supply side frictions, treat prices as continuous choices, and consider a complete information environment.

■ **Pricing with consumer switching costs.** A widely acknowledged consensus in the consumer switching cost literature is that their existence creates two countervailing effects for firms: *investing* and *harvesting* and gives them an incentive to price dynamically. Villas-Boas (2015) surveyed vast theoretical and empirical literature to conclude that depending on the assumptions about firms' and consumers' planning horizons, switching costs might result in higher (Beggs and Klemperer, 1992) or lower (Dubé et al., 2009) profits and equilibrium prices. Dubé et al. (2008) and Dubé et al. (2009) empirically analyse the implications of switching costs for monopolist's and oligopolists' pricing strategies, respectively, in the orange juice and margarine industry. The main finding of the second paper is that oligopoly profits can be lower by more than 10% in the presence of brand loyalty, since the investing motive tends to prevail. Recent contributions of Arie and Grieco (2014) and Percy (2014) establish theoretical properties of pricing equilibria under logit demand and switching costs, mostly confirming the empirical findings of Dubé et al. (2009). An important policy implication from the Arie and Grieco (2014) paper is that lower equilibrium prices resulting from brand loyalty do not necessarily translate into net welfare gains, since people might be persistently consuming lower-utility products to avoid paying the switching cost.

■ **Dynamic oligopoly with adjustment costs.** A new feature we introduce to a model with consumer-side switching cost are supply side frictions in the form of costs of adjusting prices. Early theoretical works on oligopoly models with adjustment costs include Fershtman and Kamien (1987) and Driskill and McCafferty (1989). According to the taxonomy of strategic incentives developed by Lapham and Ware (1994), Bertrand competition with price adjustment costs generates an incentive to raise prices in order to induce competitors' softer behaviour. In a duopoly, both firms turn out to be better off because of the additional costs. Jun and Vives (2003) confirm this result, noting that the nature of competition (i.e. whether it is in quantities or prices) does not really matter for the comparison of an MPE outcome with a static Nash equilibrium.

As far as empirical work in this area is concerned, Slade (1998) attempts to quantify the magnitude of fixed and variable cost of adjusting prices of saltine crackers assuming the industry is monopolistically competitive. Similarly to us, she treats the pricing decision as discrete⁸ and structurally

⁸While we explicitly say that firms can choose between two prices for each of their goods, Slade (1998) instead defines the binary decision variable, d_{it} as equal to 1 if the price goes down or up and 0 if it stays the same. Conditional on $d_{it} = 1$, the actual price is

estimates a dynamic discrete choice model. In that sense, our model extends her approach to strategic, multiproduct firms. The other important difference is that she does not model consumer behaviour at the micro level, but uses aggregate demand estimates to construct a *goodwill* state variable. Her results suggest that adjustment costs are important and vary across products and supermarket chains. In a related paper, albeit without using a dynamic structural model, Slade (1999) finds that in a strategic setting, the existence of fixed costs exacerbates price rigidities and leads to less frequent adjustments, which, if observed in many industries and aggregated, might have serious repercussions at the macro level.⁹ Kano (2013) further extends Slade's model to show that ignoring strategic interactions might lead to a biased estimate of menu costs. Recently, Ellison et al. (2015) studied dynamic pricing with adjustment costs using high-frequency data from an online seller, attributing price stickiness to managerial frictions.

■ **Econometrics of dynamic games.** Last but not least, our work broadens the scope of applications of dynamic games estimation techniques, pioneered in the seminal works of Aguirregabiria and Mira (2007), Bajari et al. (2007) and Pesendorfer and Schmidt-Dengler (2008). To the best of our knowledge, pricing problems have not been analysed in this framework to date, even though slightly different structural econometric models have been used to study dynamic pricing, see e.g. Jofre-Bonet and Pesendorfer (2003) who consider an auction setting, Sweeting (2012) whose model is single-agent and Lee, Roberts, and Sweeting (2012) for a similar problem of pricing perishable goods by a price leader and fringe competitors.

1.3 Data and industry background

■ **Data.** The data used in this chapter come from Kantar Worldpanel, which is a representative, rolling survey of UK households documenting their daily grocery purchases between November 2001 and November 2012. The average sample size for the wave starting in 2006 is around 25,000 households and for each of their shopping trips, SKUs (barcodes), prices, quantities and store of purchase are recorded at a daily frequency, together with product characteristics and indicators of promotional status.¹⁰ To find a balance between analysing a stationary environment with no new product introduction and negligibly little repositioning, and having enough variation in the data,

determined randomly according to an empirical distribution of observed price changes.

⁹Hence, our work is also loosely related to the menu cost literature, see Levy et al. (1997), Dutta et al. (1999), Guimaraes and Sheedy (2011).

¹⁰Various subsamples of this large data set have been used in previous research on consumer behaviour, such as Griffith, Leibtag, Leicester, and Nevo (2009), Seiler (2013), Dubois, Griffith, and Nevo (2014), and therefore we refer the reader to these papers for details regarding the data collection procedure.

we restrict our attention to a 200-week subsample from 2009 to 2012.

We chose to focus on the butter and margarine industry for a variety of reasons. The products involved are regularly purchased, branded and expenditures within this category make up a small part of households' budgets,¹¹ so depending on individual preferences, there is both room for brand loyalty and switching. Moreover, dairy products are perishable and have a relatively short shelf life. This suggests that stockpiling is limited and allows us to abstract from dynamic considerations on the demand side.

■ **Sales channels.** The most important sales channels for the manufacturers are the four largest supermarket chains. More than 83% of purchases recorded in our sample were made in one of the four: Asda, Morrisons, Sainsbury's or Tesco. As shown in table 1.1, their market shares are stable year-to-year and Tesco is a clear market leader. Among the big 4 chains, Morrisons has consistently the lowest market share. The fifth largest supermarket chain, Co-op, caters on average only for 3% of the market. Given the relative importance of the 4 big supermarkets in the UK market, in what follows we will focus our attention only in purchases of butter and margarine observed in Asda, Morrisons, Sainsbury's and Tesco.

Table 1.1: Expenditure shares of main supermarket chains in the butter and margarine category.

STORE OF PURCHASE	Year				
	2009	2010	2011	2012	2009-2012
Aldi	1.61%	1.61%	2.19%	3.10%	2.32%
Asda	19.52%	18.94%	19.59%	20.22%	19.58%
Co-op	2.54%	3.27%	3.19%	2.91%	3.01%
Iceland	1.85%	2.03%	2.04%	2.01%	1.99%
Lidl	2.44%	2.53%	2.58%	2.69%	2.56%
Morrisons	14.43%	14.40%	14.70%	14.35%	14.47%
Netto	1.31%	1.11%	0.49%	-	1.08%
Sainsbury's	15.18%	16.27%	15.91%	15.14%	15.64%
Tesco	34.00%	33.69%	33.66%	33.70%	33.77%
Waitrose	1.83%	1.99%	1.92%	1.88%	1.91%

Note: Shares defined as sum of expenditures on butter and margarine in a given chain during the period of interest (year) divided by total expenditures in all stores. Four biggest chains and their average market shares were highlighted. Netto sold their stores to Asda in 2011. Source: own calculations using the Kantar data.

■ **Producers.** The market is dominated by three big players: Arla, Dairy Crest and Unilever. Within each of the four retail chains, their products comprise from 75% (Tesco) to approximately 80% (Asda) of total sales. Each supermarket has also its own brand. Adding the store brand, the

¹¹The annual value of UK butter and margarine industry in 2014 is estimated to be £1.35bn.¹² Yet, at the household level, purchases of goods belonging to this category make up slightly more than 1% of total grocery expenditures (Griffith et al., 2017).

four-firm concentration ratio, CR_4 exceeds 90%.¹³ The remaining manufacturers are either small dairies that cater local markets (such as Dale Farm Dairies in Northern Ireland), or firms that are big players in other industries.¹⁴ Two of the three market leaders, Arla and Dairy Crest, are also major manufacturers of other dairy products (milk, cheese and yogurt), while Unilever is world's third-biggest consumer goods producer, who at the same time is the biggest margarine manufacturer in the world. The sales of margarine make up around 5% of Unilever's total revenue.¹⁵

■ **Products.** Butter and margarine come in different pack sizes (250g, 500g, 1kg and 2kg) and formats (block and spreadable). In our detailed data set, we observe more than 100 distinct brand-pack-format combinations produced by 12 manufacturers. Four of them are the supermarkets themselves, who sell own brand products exclusively in their outlets. Since the number of distinct brand-pack size-format combinations observed in the data is substantial we will restrict our attention to the 500g spreadable segment. We decided to focus on this subsample of all products for a number of reasons: first, this is the largest segment, comprising more than 50% of industry sales, in which butters, margarines and own brand alternatives coexist in all stores. Secondly, spreadables are much less frequently used for cooking and baking than block butters and margarines. Therefore the consumption and, consequently, interpurchase times should be relatively stable. This is important both for the discrete choice assumption in the demand model, as well as for the assumption that there are no unexpected or seasonal aggregate demand shocks in our framework. The drawback of our choice is that the outside good might also include purchases of smaller packs of the same brand, e.g. 250g packs of Lurpak or Flora, so the loyalty effect may be underestimated.¹⁶

Within the 500g segment we select six branded (the largest two of Arla, Dairy Crest and Unilever in the segment) and a composite own branded product for all four largest supermarket chains. Table 1.2 lists the name, the corresponding manufacturer and the type (butter/margarine) of the 6 selected brands. In the 2009-2012 period, all the brands mentioned in table 1.2 were long-term incumbents, some of them being present in the UK for more than 40 years. Long-run market shares

¹³In Tesco, for instance, over the 4-year period of our sample, Unilever had a share of 30.3%, Arla 23.9%, Tesco store brand 21.2% and Dairy Crest 18.3%. $CR_4 = 93.7%$ Similar calculations for Asda, Morrisons and Sainsbury's are available upon request.

¹⁴Lactalis is the manufacturer of *Président* butter, whose long-run market share is around 0.5%, but it is a much more important player in the cheese industry.

¹⁵See <http://www.bloomberg.com/news/articles/2014-12-04/unilever-plans-to-split-spreads-business-into-standalone-unit> (access on March 7, 2018).

¹⁶To check that by selecting a subset of products we do not distort the market structure, we computed expenditure- and volume-based market shares using the selected sample. Compared to the entire market, firm- and brand-level market shares in the 500g spreadable segment are quantitatively proportionate, with the only exception being Arla's higher share at the cost of lower share of the store brand. This is due to the fact that, in all 4 supermarkets, the most popular own brand products are 250g block butters. Yet, the shares of store brands remain non-negligible, and hence we believe that even after narrowing down the set of products we are still able to provide a faithful depiction of the entire industry.

Table 1.2: Manufacturers and brands of top 6 branded products.

MANUFACTURER	BRAND	BUTTER/MARGARINE
Arla	Anchor	Butter
	Lurpak	Butter
Dairy Crest	Clover	Margarine
	Country Life	Butter
Unilever	Flora	Margarine
	I Can't Believe It's Not Butter	Margarine

of the brands are relatively stable, yet one observes considerable variation at a weekly level, which we will document next. See table A.6.1 in appendix A.6 for more details on long-run market shares of all products.

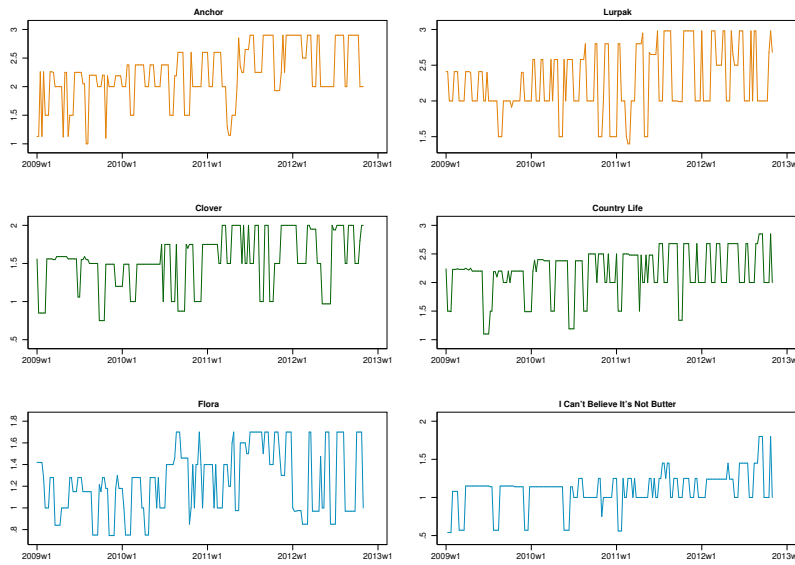
■ **Prices.** We do not have supermarket-level price data. We only observe prices actually paid by the consumers. Therefore we construct daily time series of prices for the six spreadable products listed in table 1.2 in the four big supermarket chains by taking the median price paid in a given day. This approach can be justified by the fact that after the 2000 enquiry, the Competition Commission imposed national pricing rules on the UK chain stores. This also means we do not have to impute missing prices for particular stores, because we can simply take the price observed in a different outlet of the same chain.

As with most grocery products, most price variation at the SKU level comes from periodical movements between regular (baseline) and sale prices.¹⁷ Figure 1.1 shows the evolution of prices of six 500g spreadable products in Tesco manufactured by the three biggest firms.

Usually, the regular price remains at the same level for an extended period of time, up to 18 months. For most branded products in our 200-week sample we observe a maximum of three changes of the regular price. Promotions can be as deep as 50% and the depth might vary across supermarket chains, but over 3-6 month periods one can actually observe only two price regimes for each product. As opposed to the *high-low* pricing of national brands, supermarkets employ *everyday low price* strategies for their private labels. This implies that average prices of store brands are consistently much lower than the prices of branded products – see table A.6.2 in the appendix. Within segments of the market defined by size-format combination, promotional prices of branded butters and margarines sometimes tend to match the prices of own brand products and very rarely fall below that level.


¹⁷See Hosken and Reiffen (2004) and Nakamura and Steinsson (2008) for reviews of empirical regularities about prices.

Figure 1.1: Prices of 500g spreadable butters and margarines in Tesco stores.



Note: Prices in Tesco stores between 01/01/2009 and 28/10/2012.

Figure 1.2: Price promotions and market shares of Country Life (500g spreadable) in Tesco.



Note: Normalised prices and market shares in Tesco between 01/01/2009 and 28/10/2012.

We also document that promotions have important effects on market shares. This point is illustrated in figure 1.2. In the upper part of the figure we plotted the time series of normalised prices of DC's Country Life 500g in Tesco stores. To build this time series we attributed to the product the average regular price if Country Life 500g was not marked with a promotion flag in that period and the average promotional price otherwise. In the lower part of the figure we

plotted the market shares of Country Life 500g for the same time span. All promotional periods are associated with spikes in market shares. After the promotion shares appear to return promptly to their pre-promotional levels.

In summary, the butter and margarine industry is a typical example of multiproduct oligopoly. The market is dominated by a small set of firms selling products under different brand names. Prices of these products are far from being continuous. For branded products we observe a finite and relative small number of prices during our 200 weeks sample. Most of the price changes are between regular and sales prices and price promotions have a clear effect on market shares. Store brands are also important in the industry. Prices of spreadable products sold under store brands are more stable and usually lower than promotional prices of branded products. These elements will play a prominent role in the construction of our dynamic pricing model. The two other building blocks of our model: price adjustment costs and consumer loyalty, will be discussed in detail in the next section.

1.4 Descriptive evidence

This section provides preliminary evidence of consumer loyalty and price inertia in the context of the UK butter and margarine industry. We document brand-switching patterns at the individual level and look at the persistence and rigidity of retail prices. We conclude with a set of reduced-form regressions which shed light on the relationship between current pricing decisions, past prices and past market shares.

1.4.1 Consumer loyalty

Consumers that are loyal to a brand are expected to buy the same brand more often. Loyalty, in other words, implies that consumer choices exhibit some degree of inertia over time. In order to check whether this behaviour is present in our data, we investigated brand switching patterns observed in the household-level data.

For all purchases recorded in the data set, we calculated the empirical frequencies corresponding to the transition probabilities of the loyalty state, according to two different definitions of loyalty. Our first definition of consumer loyalty (see table 1.3) implies that consumers' memory reaches only one period back – see Horsky, Pavlidis, and Song (2012) and Eizenberg and Salvo (2015) – i.e., according to the first definition, loyalty will be calculated as the fraction of households choosing the same brand between *two adjacent weeks*. Our second definition is consistent

with Dubé, Hitsch, Rossi, and Vitorino (2008), Dubé, Hitsch, and Rossi (2009) and Pavlidis and Ellickson (2017). These papers define consumer loyalty as the fractions of consumers that buy the same brand between *two subsequent purchases* and, therefore, the time between purchases does not matter (table 1.4).

Table 1.3: Consumer switching patterns for purchases made in two subsequent weeks.

Purchase at t	Purchase at $t + 1$						
	ANC	LUR	CLO	COU	FLO	ICB	SB
ANCHOR	73.59%	6.62%	2.71%	5.21%	6.34%	2.56%	2.97%
LURPAK	3.64%	80.27%	1.52%	3.52%	5.43%	2.21%	3.41%
CLOVER	2.05%	2.93%	73.83%	1.97%	8.92%	5.48%	4.82%
COUNTRY LIFE	6.74%	9.18%	3.27%	68.40%	5.61%	3.19%	3.61%
FLORA	2.55%	4.24%	4.39%	1.79%	75.08%	6.58%	5.37%
ICBINB	1.71%	2.70%	3.85%	1.44%	10.37%	72.14%	7.79%
STORE BRAND	2.13%	4.79%	4.15%	1.80%	8.93%	8.99%	69.20%

Note: Frequencies based on a sample of 126,508 individual purchases between 01/2009 and 10/2012. Store brand here is a composite generic good including Asda, Morrisons, Sainsbury's and Tesco own brand products. The highlighted entries on the diagonal denote the percentage of loyalty-driven purchases.

Table 1.4: Consumer switching patterns

Purchase at t	Subsequent purchase						
	ANC	LUR	CLO	COU	FLO	ICB	SB
ANCHOR	62.82%	9.22%	4.02%	7.22%	8.23%	3.92%	4.57%
LURPAK	4.49%	74.26%	2.21%	4.31%	6.98%	3.09%	4.66%
CLOVER	2.83%	3.63%	58.98%	2.42%	14.59%	10.05%	7.50%
COUNTRY LIFE	9.71%	12.51%	4.88%	54.25%	8.02%	4.83%	5.80%
FLORA	2.80%	4.72%	6.19%	1.96%	65.47%	10.36%	8.50%
ICBINB	2.12%	3.26%	6.31%	1.72%	17.08%	56.90%	12.61%
STORE BRAND	2.16%	4.70%	5.17%	2.01%	11.93%	12.15%	61.89%

Note: Frequencies based on a sample of 569,338 individual purchases between 01/2009 and 10/2012. Store brand here is a composite generic good including Asda, Morrisons, Sainsbury's and Tesco own brand products. The highlighted entries on the diagonal denote the percentage of loyalty-driven purchases.

The first striking observation about the two tables is that, regardless of the definition, loyalty seems to play a decisive role in the determination of consumer choices. Restricting our attention to purchases in the two subsequent weeks, we observe a stronger loyalty effect, which might indicate that some consumers exhibit shorter memory.¹⁸ The fractions of loyalty-driven purchases are relatively similar across products. Even though brand commitment seems to play a key role in this industry, there is still a fair number of consumers who switch products every period and firms

¹⁸Since 22.2% of observations used to calculate the transition probabilities in table 1.4 are the same ones as the data used to construct table 1.3, we checked what part of the loyalty effect wears off after 1 week. When the interpurchase time is more than 2 weeks, the fractions are approximately 10 p.p. lower than the ones in table 1.4. After 10 weeks, about 40% of consumers are still loyal to national brands, whereas the effect for store brand disappears almost completely.

might be willing to price aggressively to fight for them. In line with the intuition, consumers who bought butter last period will be more willing to buy butter again next time, rather than switch to margarine or store brand. Rather not surprisingly, switching to store brand is especially popular among consumers of the cheapest margarine brand – *I Can't Believe It's Not Butter*.

1.4.2 Price rigidities

We look at the frequency of price changes and duration of promotions as possible indicators for price adjustment costs for the suppliers. In table 1.5 we provide descriptive statistics for weekly price changes based on the 200-week period. Given our selection of products, the maximum number of changes is 2 for each of the firms. The means in table 1.5 reveal that prices, on average,

Table 1.5: Frequency of price changes.

FIRM	MEAN	STD. DEV.	%(1)	%(2)
<i>ASDA</i>				
Arla	0.347	0.527	29.65%	2.51%
Dairy Crest	0.357	0.521	31.66%	2.01%
Unilever	0.271	0.493	22.65%	2.21%
<i>MORRISONS</i>				
Arla	0.412	0.560	34.17%	3.52%
Dairy Crest	0.342	0.545	27.14%	2.01%
Unilever	0.277	0.461	26.55%	0.56%
<i>SAINSBURY'S</i>				
Arla	0.472	0.687	25.13%	11.06%
Dairy Crest	0.317	0.591	18.59%	6.53%
Unilever	0.281	0.483	25.13%	1.51%
<i>TESCO</i>				
Arla	0.533	0.695	30.15%	11.56%
Dairy Crest	0.437	0.631	28.64%	7.54%
Unilever	0.469	0.673	26.77%	10.10%

Note: Table presents average number of per-firm weekly price changes (without specifying direction) in each of the supermarket chains. Fourth and fifth column show the percentage of weeks with 1 and 2 price changes, respectively.

change much less frequently than every period. For all firms, adjustments occur most often in Tesco, with both three firms having approximately 1 price change every other week. In the remaining three retailing chains, Unilever is the least likely to change its prices – 75% of the time it makes no adjustments, while Arla changes prices of both *Anchor* and *Lurpak* in Sainsbury's and Tesco on more than 10% of all occasions.

Naturally, the full picture is much more complicated,¹⁹ than what we see in table 1.5 but even

¹⁹For simplicity, we not only abstracted from all strategic aspects of price adjustments here, but also did not specify which

at this very general level we can still derive some implications relevant to our structural model. First, we may expect the adjustment costs to be non-negligible for all firms. Our second hypothesis is that they vary across firms and, to a lesser extent, across supermarkets, just like in Slade (1998). Third, the descriptives show no evidence of any form of economies of scope, or synergies, since adjustments of more than one price per firm occur relatively rarely.

Another piece of evidence suggesting that costs of adjusting prices might play an important role in this industry is presented in table 1.6.²⁰ For all six products we observe between 20 and 27 distinct sale spells in the 200-week sample. Average duration of a single spell is around 3-4 weeks, depending on the brand. However, we also witness much shorter and much longer periods of reduced prices, varying from one to as many as 20 weeks. These numbers seem to be an additional piece of evidence for the existence of price rigidities, and because of relatively high dispersion in the duration of sales, we can also exclude the possibility that promotions always last for a fixed number of weeks. Within the context of our model, we would be observing longer sale spells if a firm fails to attract sufficient number of new consumers immediately after decreasing the price.

Table 1.6: Durations of promotions in Tesco stores.

PRODUCTS BY MANUFACTURER	# SPELLS	MEAN DUR.	STD. DEV.	MIN	MAX
Arla					
ANCHOR	27	3.74	2.11	1	10
LURPAK	27	4.52	3.58	1	20
Dairy Crest					
CLOVER	20	3.80	2.04	1	9
COUNTRY LIFE	24	3.12	1.39	1	6
Unilever					
FLORA	26	3.96	3.54	1	16
ICBINB	22	3.54	2.11	1	9

Note: # SPELLS denotes the number of distinct promotional spells in the 200-week sample. Remaining columns describe the distribution of durations of sales.

1.4.3 Implications of consumer loyalty and price adjustment costs for price dynamics

The evidence shown so far in this section suggests there is a high degree of inertia in consumer choices and equilibrium prices. Inertia in choices may be directly related to consumer loyalty, while inertia in prices can be attributed to the presence of price adjustment costs in the industry.

movements are upward and which downward.

²⁰For the sake of brevity, we only present results from Tesco here, see appendix A.6 for the remaining results.

We now examine how these two elements affect price decisions.

To do this we estimate a series of descriptive regressions of current prices on past prices and shares. Intuitively, if price adjustment costs are relevant then, all else constant, brand j 's current price will be correlated with its past price. Analogously, one way of uncovering the importance of consumer inertia to price dynamics is through the analysis of (partial) correlations between brand j 's current price and its past share. Inertia in consumer choices imply that, all else constant, market share of brand j at time t will depend positively on its market share in the past. This dependence, in turn, may change manufacturers' incentives to set prices in the present period.

The set of observed prices for the products in our sample is relatively small. In particular, price movements generally correspond to switches between regular and promotional prices. For this reason, instead of modelling prices directly we model promotion decisions, i.e. the dependent variable in our models will be a dummy variable that assumes 1 if brand j was marked with a promotional flag in a given period and zero otherwise.²¹ More specifically the regression equation that we study takes the form:

$$a_{jmt} = \alpha + \sum_k \beta_k \cdot a_{km,t-1} + \sum_k \gamma_k \cdot s_{km,t-1} + \alpha_{jm} + \epsilon_{jmt}, \quad [1.1]$$

where, a_{jmt} is a dummy variable that assumes 1 if brand j was in promotion in supermarket m in period t and zero otherwise; s_{jmt} is the market share of brand i in supermarket m in period t ; α_{jm} is a supermarket/brand fixed-effect; ϵ_{jmt} is an idiosyncratic error term and $(\alpha; \{\beta_k\}_k; \{\gamma_k\}_k)$ are coefficients to be estimated. The second summation includes (past) shares of all branded products plus the share of the store brand. The first summation only includes (past) actions of branded products. We do not include the store brands because there is little variation in their prices (see section 1.3) and their effects will not be identifiable with a fixed effects regression.

We present the results of this regression in table 1.7. To estimate the coefficients of interest we stacked data for Asda, Morrisons, Sainsbury's and Tesco and used a fixed effects estimator.²² Each column has the estimated coefficients for each of the 6 brands. We emphasise two important findings. First, the coefficients in the main diagonal of the upper part of the table show that current actions depend positively on past actions. The coefficients are large and significant at 1%. In particular, the point estimates imply that if a brand had been sold at promotional prices during the previous week it is 46-58 percentage points (depending on the brand) more likely to be sold at

²¹ As it will become clear later on, this formulation is also consistent with some practical assumptions that we have to make in order to estimate and solve the structural model.

²² We also run OLS regressions separately for each supermarket. The results of these regressions are qualitatively and quantitatively similar to those in table 1.7. For brevity they are not shown.

Table 1.7: Price regressions

	Arla		DC		Unilever	
	ANCHOR	LURPAK	CLOVER	C. LIFE	FLORA	ICB
\mathbf{a}_{t-1}						
Price Anchor	0.465*** (0.02)	-0.074 (0.05)	0.008 (0.01)	-0.022 (0.04)	-0.016 (0.01)	0.026 (0.02)
Price Lurpak	-0.058 (0.05)	0.513*** (0.06)	0.009 (0.02)	-0.059 (0.05)	-0.025 (0.05)	0.043 (0.06)
Price Clover	-0.025 (0.05)	-0.044 (0.04)	0.547*** (0.03)	0.019 (0.01)	-0.000 (0.01)	-0.015 (0.03)
Price Country Life	-0.000 (0.04)	-0.012 (0.01)	-0.011 (0.02)	0.534*** (0.04)	0.001 (0.01)	-0.045 (0.03)
Price Flora	-0.027 (0.03)	0.025 (0.01)	-0.048 (0.03)	0.008 (0.04)	0.523*** (0.05)	-0.073* (0.03)
Price ICBINB	-0.032 (0.03)	0.025 (0.03)	0.000 (0.02)	-0.096** (0.03)	0.019 (0.03)	0.585*** (0.04)
\mathbf{s}_{t-1}						
Share Anchor	3.523 (1.51)	1.680 (1.73)	1.332 (2.15)	-0.058 (1.69)	1.169 (1.34)	1.374 (1.80)
Share Lurpak	0.774 (1.22)	-0.113 (0.98)	-1.171 (0.76)	1.805 (1.36)	0.298 (1.07)	0.505 (0.53)
Share Clover	0.185 (0.11)	1.270 (0.60)	2.069*** (0.32)	-0.322 (0.49)	-0.336 (0.28)	0.348 (0.33)
Share Country Life	-2.750** (0.69)	0.659 (1.68)	1.817 (2.11)	3.308 (2.54)	-1.764 (2.53)	2.657 (1.69)
Share Flora	0.295 (0.37)	0.289 (0.42)	0.085 (0.41)	-0.874** (0.26)	1.226* (0.52)	0.105 (0.45)
Share ICBINB	0.572 (0.37)	-0.442 (0.49)	-0.317 (0.21)	1.320*** (0.19)	0.417 (0.42)	1.723*** (0.22)
Share Store Brand	-1.262 (1.34)	-0.689 (0.87)	-0.198 (1.35)	-1.189 (1.01)	-0.387 (0.48)	-1.947* (0.65)
Observations	796	796	796	796	796	796
R-squared	0.289	0.298	0.393	0.346	0.335	0.488

Note: fixed effects regressions for each brand for 4 cross-sections (supermarkets). \mathbf{a}_{t-1} refer to promotional status in the previous period and \mathbf{s}_{t-1} to shares in the previous period. Significance levels: *** 1%, ** 5%, * 10%.

promotional prices during the current week. In line with the patterns shown in tables 1.5 and 1.6, these results indicate that sources of price rigidities such as adjustment costs can play an important role in this industry. Second, the effect of consumer loyalty on price decisions does not seem to be as important as the effect of price adjustment costs on current prices – in spite of the strong correlations between past and current consumer choices (see tables 1.3 and 1.4). Coefficients attached to (own) past shares are significant at 10% for 3 of the 6 brands only. Furthermore, when these coefficients are significant, their magnitudes appear to be small. For example, an increase of 1 percentage point in the market share of Clover at $t - 1$ increases the probability of a Clover promotion in 2 percentage points at period t . It is worth noting that the off-diagonal coefficients suggest that the role of strategic interactions between firms may be limited, but it cannot be ruled out completely. Our structural analysis below will proceed under the assumption that firms are

competing in a pricing game, which encompasses a collection of single-agent decision problems.

It is worth emphasising that the regressions in this section do not take into account potentially important features of the industry, such as multiproduct nature of the firms and forward-looking behaviour. Moreover, the fact that the coefficients corresponding to lagged market shares are low can be explained by the fact that in some states of the world firms prefer to invest, which would imply positive correlation between the promotional status dummy and lagged market share, and in other states their profit-maximising decision is to harvest, implying negative correlation. A simple linear specification cannot disentangle these mutually offsetting effects. All those different forces can however be captured by our structural model studied in the following section.

1.5 Model

The descriptive analysis in the previous section suggests that price rigidities and consumer inertia are important in the UK butter and margarine industry. This section develops a dynamic model to explain how pricing decisions in the butter and margarine industry are affected by these two elements. Since we treat pricing decisions as discrete choices, the way we describe the game draws heavily on concepts developed in the empirical dynamic discrete games literature.²³ We characterise the equilibrium of the model and discuss the identification of its primitives below.

The model we present involves forward-looking, multiproduct oligopolists engaged in a dynamic pricing game over non-durable goods. The consumers in our model are assumed to be myopic, so their expectations about future prices do not play any role in their contemporaneous choices.²⁴ Instead, dynamic pricing incentives arise from brand loyalty, which can be alternatively interpreted as inertia or switching cost. Firms offer temporarily lower prices to attract new or returning consumers and use the fact that at least some of them will remain loyal when the price returns to the regular (high) level.

The second dynamic component of our model is the cost incurred when prices are being adjusted, so that the framework falls into the wider class of oligopoly games with adjustment costs,²⁵ and can be seen as an extension of Slade's (1998) single-agent, one-product model to a multiprod-

²³For more details on this class of models, we refer the reader to recent literature surveys, for example Arcidiacono and Ellickson (2011) or Aguirregabiria and Nevo (2013).

²⁴Apart from tractability, the problem with having both forward-looking firms and consumers is that in a Markov perfect equilibrium the information sets of consumers and firms are identical, and so are their expectations regarding future states of the world, as in Goettler and Gordon (2011). In order to avoid making this assumption, one should change the equilibrium notion to explicitly allow for asymmetric information between firms and consumers, for example by adapting the framework of Fershtman and Pakes (2012).

²⁵See e.g. Fershtman and Kamien (1987), Lapham and Ware (1994), Jun and Vives (2003), and chapter 9 of Vives (2002) for a broader perspective.

uct setting with strategic players.²⁶

1.5.1 Preliminaries

We consider a discrete time game with infinite horizon, where periods are denoted by $t = 1, \dots, \infty$. Firms, indexed by $i = 1, \dots, N$, compete over a discounted sum of profits in a single market,²⁷ by choosing whether to charge low or high price for each of the goods they produce, where the low/high prices can vary across products.²⁸ The set of products offered by firm i is $\mathcal{J}_i = \{i_1, i_2, \dots, i_{|\mathcal{J}_i|}\}$, where $|\cdot|$ denotes the cardinality of a set. We let all products be differentiated, so that the entire set of products available to the consumer is $\mathcal{J} = \bigcup_{i=1}^N \mathcal{J}_i$. We assume the market is mature and rule out entry of firms and introduction of new products.²⁹

On the demand side, there is a mass H of households. We assume they are myopic and face a discrete choice problem in each period. By assumption H does not change over time. The same households visit the stores each period, but they have the option of choosing the outside good.

The sequence of events within each period is as follows. First, firms observe last period's prices, demand realisations and a random draw from the distribution of private cost shocks. Based on this information, they simultaneously choose between regular (high) or promotional (low) prices for all products they manufacture. If the prices differ from last period's ones, they pay an adjustment cost.³⁰ After the prices are set, consumers make purchases, firms learn the realisation of demand and receive period profits. The game moves on to the next period and state variables update according to their transition laws.

1.5.2 Firms

Let \mathcal{A}_i denote the set of actions available to player i . Since, by assumption, there are two regimes for the price of each good and the prices are set simultaneously for the entire portfolio of products of each firm, this is a finite set with cardinality equal to $|\mathcal{A}_i| = 2^{|\mathcal{J}_i|}$. For example, if $|\mathcal{J}_i| = 2$, player i can choose among 4 actions and $\mathcal{A}_i = \{(p_{i_1}^H, p_{i_2}^H), (p_{i_1}^H, p_{i_2}^L), (p_{i_1}^L, p_{i_2}^H), (p_{i_1}^L, p_{i_2}^L)\}$, where p^H and p^L denote high and low price, respectively.

In general, there is no need to assume that the number of possible actions is the same for all

²⁶Building on Slade (1998), Slade (1999) and Kano (2013) emphasise why strategic interactions might matter for the estimates of adjustment costs.

²⁷In the empirical application of the model, we define market as a single national retailing outlet.

²⁸It is straightforward to extend the model to allow for the pricing decision to be choice between more than two possible values. As long as the decision is discrete (number of actions finite), our results hold.

²⁹In section 1.3 we showed that this assumption can be plausibly maintained in our data.

³⁰We relegate a detailed interpretation of this cost to section 1.6.

players, or make any other symmetry assumptions. Note also that for single-product firms, the action space is binary and the structure of the problem becomes very similar to an entry/exit game.

The problem of firm i in period t is to choose an action $a_{it} \in \mathcal{A}_i$ to maximise the expected discounted stream of payoffs: $\mathbb{E}_t \sum_{\tau=t}^{\infty} [\beta^{\tau-t} \Pi_i(\mathbf{a}_{\tau}, \mathbf{z}_{\tau}, \varepsilon_{i\tau}(a_{i\tau}))]$, where $\beta \in (0, 1)$ is the discount factor and $\Pi_i(\cdot)$ denotes firm i 's profit in period t . $\mathbf{a}_t = (a_{1t}, a_{2t}, \dots, a_{Nt})$ collects the actions of all players. Occasionally we will abuse the notation and write $\mathbf{a}_t = (a_{it}, \mathbf{a}_{-it})$. $\mathbf{z}_t \in \mathcal{Z}$ is the vector of publicly observed, payoff-relevant state variables, which in our model contains last period's market shares and actions, so $\mathbf{z}_t = (\mathbf{s}_{t-1}, \mathbf{a}_{t-1})$, and $\varepsilon_{it} = (\varepsilon_{it}(a_i))_{a_i \in \mathcal{A}_i}$ is a vector of iid private cost shocks associated with a_i . The expectation is taken over the distribution of beliefs regarding other players' actions, next period's draws of ε , and the future evolution of state variables.

We assume that the private shock enters the profit function additively, so that it can be expressed as:

$$\begin{aligned} \Pi_i(\mathbf{a}_t, \mathbf{z}_t, \varepsilon_{it}) &= \pi_i(a_{it}, \mathbf{a}_{-it}, \mathbf{s}_{t-1}) + \sum_{\ell \in \mathcal{A}_i} \varepsilon_{it}(\ell) \cdot \mathbf{1}(a_{it} = \ell) \\ &+ \sum_{\ell \in \mathcal{A}_i} \sum_{\ell' \neq \ell} SC_i^{\ell' \rightarrow \ell} \cdot \mathbf{1}(a_{it} = \ell, a_{i,t-1} = \ell'), \end{aligned} \quad [1.2]$$

where $SC_i^{\ell' \rightarrow \ell}$ is the adjustment cost of switching from action ℓ' to ℓ and $\mathbf{1}(\cdot)$ is the indicator function. The first part of (1.2) is the static profit accrued over the time period and can be written as:

$$\pi_i(a_{it}, \mathbf{a}_{-it}, \mathbf{s}_{t-1}) = H \cdot \sum_{j \in \mathcal{J}_i} (p_{jt}(a_{it}) - mc_j) \cdot s_{jt}(a_{it}, \mathbf{a}_{-it}, \mathbf{s}_{t-1}) \quad [1.3]$$

We use the notation $p_{jt}(a_{it})$ to emphasise the 1-to-1 relationship between firm's action and the price of product j . mc_j is a constant marginal cost and s_{jt} is the market share derived from the consumer's problem. To keep the notation parsimonious, fixed cost of operating is normalised to zero, although one can still interpret the ε 's as shocks shifting fixed costs from period to period. As we do not consider firm entry or exit, this is a fairly innocuous assumption.³¹

Rewriting the expectation in terms of beliefs and perceived transition laws, firm i 's best re-

³¹In principle, we could still have fixed cost in the profit function while describing the theoretical model. However, in contrast to the entry and exit game, it would not matter for the optimal choice of strategy, since it would appear on both sides of all the inequalities defining firms' best responses. From an econometric point of view, this structural parameter would not be identified (see Aguirregabiria and Suzuki (2014) for an extensive discussion of this problem pertaining to a dynamic entry game).

sponse is a solution to the following Bellman equation:

$$V_i(\mathbf{z}_t, \varepsilon_{it}) = \max_{a_{it} \in \mathcal{A}_i} \left\{ \sum_{\substack{\mathbf{a}_{-it} \in \times_{j \neq i} \mathcal{A}_j}} \sigma_i(\mathbf{a}_{-it} | \mathbf{z}_t) \cdot [\Pi_i(a_{it}, \mathbf{a}_{-it}, \mathbf{z}_t, \varepsilon_{it}) + \beta \sum_{\mathbf{z}_{t+1} \in \mathcal{Z}} G(\mathbf{z}_{t+1} | \mathbf{z}_t, \mathbf{a}_t) \int V_i(\mathbf{z}_{t+1}, \varepsilon_{it+1}) dQ(\varepsilon_{it+1})] \right\} \quad [1.4]$$

In the expression above, we used the notation $\sigma_i(\mathbf{a}_{-it} | \mathbf{z}_t)$ to denote firm i 's beliefs that given the state variable realisation \mathbf{z}_t , its rivals will play an action profile \mathbf{a}_{-it} . If $\mathbf{a}_{-it} = (\ell_1, \dots, \ell_{i-1}, \ell_{i+1}, \dots, \ell_N)$, by independence of private information in equilibrium we have:

$$\sigma_i(\mathbf{a}_{-it} | \mathbf{z}_t) = \prod_{k \neq i} \Pr(a_{kt} = \ell_k | \mathbf{z}_t)^{\mathbf{1}(a_{kt} = \ell_k)} \quad [1.5]$$

where $\Pr(\cdot)$ is the conditional choice probability (CCP, Hotz and Miller (1993)). In the second part of expression (1.4), we implicitly assumed that the joint transition probabilities of public and private state variables are conditionally independent and can be factorised as $G(\mathbf{z}_{t+1} | \mathbf{z}_t, \mathbf{a}_t)Q(\varepsilon_{it+1})$. This is a standard practice in the dynamic games literature (see assumption 2 in Aguirregabiria and Mira (2007) or M2 in Sanches et al. (2016b)).

Before we define the equilibrium, we will outline the consumer's problem in our model. In section 1.6 we show how we recover the primitives of the game from the observed data, which, according to the notation we adopted here, are $\{\pi_i, \mathbf{SC}_i, G, \beta, Q\}_{i=1}^N$, where \mathbf{SC} is a vector of adjustment costs.

1.5.3 Consumers

Consumers in our setting are assumed to arrive in the market every week and choose one product $j \in \mathcal{J}$ or refrain from buying anything, thus picking the outside option of not buying anything, which, following the usual convention, we denote using the subscript 0. As mentioned earlier, in contrast to firms, individuals in our model are not sophisticated and their decisions are myopic. However, at the instant of purchase, they still remember what their previous choice was, as it directly affects their current utility. We consider one possible interpretation of state dependence, namely product (brand) loyalty.³² For loyal consumers, current utility is higher if they purchase the same product they did on the previous occasion, and firms have an incentive to charge temporarily lower prices in order to build up a base of loyal customers who will be willing to pay a higher

³²Dubé, Hitsch, and Rossi (2010) examine other explanations for state dependence in discrete choice models of demand for margarine and orange juice, such as search and learning. Their main finding shows that only loyalty is not rejected in the scanner data set they use.

price in the future. The presence of an outside good allows us to account for the fact that not all consumers make purchases every week, while we remain agnostic about their consumption habits.

Importantly, the complexity of the aggregate demand function derived from the individual-level problem should be a compromise between tractability and realism. Moreover, it has to satisfy the Markov property, as seen in (1.3). Already because of the fact that consumers are not forward-looking, we avoid having to model their expectations about future prices of all products. A richer specifications of consumer heterogeneity can lead to a sharp increase in the dimension of the state space,³³ but on the other hand, some theoretical models of sales and dynamic pricing (e.g. Conlisk et al. (1984), Hendel and Nevo (2013)) emphasise that sales arise as a result of price discrimination between groups of consumers with different preferences.

For the purpose of this study, we assume out persistent consumer heterogeneity. This stands in contrast to Dubé et al. (2009) and Pavlidis and Ellickson (2017), who both allow for random coefficients in their demand models, but are forced to make arbitrary simplifying assumptions while analysing the supply side game to limit the dimension of the state space and keep the problem tractable.

In what follows, we assume that individual household, indexed by h , chooses an alternative from the set $\mathcal{J} \cup \{0\}$ to maximise its contemporaneous utility given by:

$$u_{jt}^h = \delta_j - \eta \cdot p_{jt} + \gamma \cdot \mathbf{1}(y_{t-1}^h = j) + \xi_{jt}^h \quad j = 1, \dots, |\mathcal{J}| \quad [1.6]$$

$$u_{0t}^h = \xi_{0t}^h \quad [1.7]$$

δ_j are alternative-specific intercepts, fixed over time. $\mathbf{1}(y_{t-1}^h = j)$ equals one if household h 's purchase at $t - 1$ was the same as the one in the current period. γ is a parameter measuring the degree of consumer loyalty (if $\gamma > 0$).³⁴ In this setting, there is no persistent consumer heterogeneity and households differ only by the realisation of their private shocks, ξ_{jt}^h , and their previous purchase.

Under the assumption that ξ 's are independent type-I extreme value shocks, the probability that household h will purchase good j at time t is:

$$\Pr_t^h(j | \mathbf{p}_t, y_{t-1}^h) = \frac{\exp(\delta_j - \eta \cdot p_{jt} + \gamma \cdot \mathbf{1}(y_{t-1}^h = j))}{1 + \sum_{g=1}^{|\mathcal{J}|} \exp(\delta_g - \eta \cdot p_{gt} + \gamma \cdot \mathbf{1}(y_{t-1}^h = g))} \quad [1.8]$$

³³More specifically, since lagged market shares are a part of \mathbf{z}_t , with H consumer types we would have $H \cdot |\mathcal{J}| + N$ payoff-relevant state variables to keep track of.

³⁴In principle it is possible to have one loyalty parameter for each good in the choice set. From the descriptive evidence in tables 1.3 and 1.4 we do not see, however, clear differences in the patterns of brand choice inertia. To keep the model parsimonious we assume that the loyalty parameter is the same across brands.

Note that, according to (1.8), the only dimension of household heterogeneity is reflected in purchase history. Since we are ultimately interested in aggregate market shares, we can use the law of total probability to integrate it out (omitting conditioning sets and superscripts to ease notation):

$$\Pr_t(j) = \sum_{g=0}^{|\mathcal{J}|} \Pr_{t-1}(g) \cdot \Pr_t(j | \mathbf{p}_t, y_{t-1} = g) \quad [1.9]$$

What we call $\Pr_t(j)$ in (1.9) is the same as s_{jt} in (1.3), just like in the standard multinomial logit model. Since characteristics of the goods do not change over time, we can remove them from the set of payoff-relevant state variables, and therefore aggregate market shares are characterised by the following Markov process (Horsky et al., 2012):

$$\begin{aligned} s_{jt}(a_{it}, \mathbf{a}_{-it}, \mathbf{s}_{t-1}) &= \sum_{g=0}^{|\mathcal{J}|} s_{g,t-1} \cdot \Pr_t(j | \mathbf{p}_t(\mathbf{a}_t), y_{t-1} = g) & [1.10] \\ &= s_{0,t-1} \frac{\exp(\delta_j - \eta \cdot p_{jt})}{1 + \sum_{g=1}^{|\mathcal{J}|} \exp(\delta_g - \eta \cdot p_{gt})} \\ &\quad + s_{j,t-1} \frac{\exp(\delta_j - \eta \cdot p_{jt} + \gamma)}{1 + \sum_{g=1}^{|\mathcal{J}|} \exp(\delta_g - \eta \cdot p_{gt} + \gamma \cdot \mathbf{1}(g = j))} \\ &\quad + \sum_{\substack{g=1 \\ g \neq j}}^{|\mathcal{J}|} s_{g,t-1} \frac{\exp(\delta_j - \eta \cdot p_{jt})}{1 + \sum_{g'=1}^{|\mathcal{J}|} \exp(\delta_{g'} - \eta \cdot p_{g't} + \gamma \cdot \mathbf{1}(g' = g))} \end{aligned}$$

Since $\sum_{g=0}^{|\mathcal{J}|} s_{g,t} = 1$ for all t , (1.10) can be further rearranged so that firms do not have to keep track of an additional state variable (share of “no purchases” every period).

In our baseline specification, we assume that consumers’ memory reaches only one period back. This approach, suggested by Horsky, Pavlidis, and Song (2012) and recently employed by Eizenberg and Salvo (2015), might not be the optimal way of modelling consumer loyalty, but is the only one in which, after aggregation, firms can keep track of past market shares to predict current demand. Dubé, Hitsch, Rossi, and Vitorino (2008), Dubé, Hitsch, and Rossi (2009), and Pavlidis and Ellickson (2017), on the other hand, define the state variables as the fractions of consumers loyal to each of the goods at the beginning of each period. We present this alternative specification in appendix A.1. Conceptually, the key difference at the micro level is that choosing the outside option does not change the loyalty state of an individual consumer. As appealing as this sounds, there are two potential disadvantages of this specification: first, it assumes that consumers who purchase very infrequently are endowed with the same degree of loyalty as people who buy every period; secondly, looking from the firms’ perspective, the interpretation of the state variables becomes problematic.

1.5.4 Equilibrium

We focus on stationary pure strategy Markov perfect equilibria. Stationarity means that we can abstract from calendar time and omit the time subscript and assume firms will always play the same strategies upon observing the same realisation of $(\mathbf{z}, \varepsilon)$. Formally the equilibrium to this game is a vector of firms' optimal price decisions – i.e. firms solve problem (1.4) taking as given their (rational) beliefs on the actions of other players – for every possible realisation of the state vector, $(\mathbf{z}, \varepsilon)$. Since the game can be seen as a particular reinterpretation of the Ericson and Pakes (1995) dynamic oligopoly framework, the proof of equilibrium existence follows from Aguirregabiria and Mira (2007), Pesendorfer and Schmidt-Dengler (2008) and Doraszelski and Satterthwaite (2010). We refer the readers to these papers for a more detailed discussion of this equilibrium and proofs of its existence.

1.6 Identification and estimation

This section discusses the identification and the estimation of the primitives of the dynamic pricing model. We first describe the structural and parametric restrictions that we use to identify the model. Our identification is constructive and follows from a recent paper by Komarova et al. (forthcoming). We then proceed to a description of our estimation strategy. In practice the estimation procedure relies on the two step approach pioneered by Hotz and Miller (1993) with parametric approximations of the value function as in Sweeting (2013). In the last part of this section we report and discuss the structural estimates.

1.6.1 Identification

The primitives in our model are $\{\pi_i, \mathbf{SC}_i, G, \beta, Q\}_{i=1}^N$. We can break down our set of assumptions into two separate groups. In the first, we impose restrictions on the data in order to recover them stems from economic theory that involves: the timing of the game, the equilibrium concept, and the specification of the period payoff function.³⁵ In the second, we impose parametric restrictions that make our problem tractable. They include the assumptions on the distribution of demand and cost shocks, which are both assumed type-I extreme value. The demand system and the process that governs the transition of states, denoted by G , can be identified outside the dynamic programming model from data on consumer choices. Subsequently, the remaining parameters in the model to be identified are $\{\pi_i, \mathbf{SC}_i, \beta\}_{i=1}^N$. The identification of these parameters is based on the Markov

³⁵Here we are only alluding to the fact that $\pi_i(a_{it}, \mathbf{a}_{-it}, \mathbf{s}_{t-1}) = H \cdot \sum_{j \in \mathcal{J}_i} (p_{jt}(a_{it}) - mc_j) \cdot s_{jt}(a_{it}, \mathbf{a}_{-it}, \mathbf{s}_{t-1})$, and not discussing how s_{jt} depends on the prices and past market shares which relies on a parametric assumption

Perfect Equilibrium of our dynamic model. Next we discuss in details the identification of these objects.

The identification of adjustment cost parameters in our model relies on Komarova, Sanches, Silva Jr., and Srisuma (forthcoming), who show differences in adjustment cost parameters can be identified and individual adjustment costs can be identified under a normalisation.³⁶ Here we assume the producers pay an adjustment cost only when the regular price is reduced.³⁷ This assumption is motivated by the notion that price reductions are associated with different forms of promotional activities – including relocation of products to shelves with more visibility, leaflets printing, production of TV commercials, compensating for retailers’ lower markups etc. – and that these costs are ultimately paid by the manufacturers. The existence of such fees has been documented both in the marketing literature (cf. Kadiyali et al. (2000), Chintagunta (2002)), popular press and industry reports which reveal that *“70% of supermarket suppliers make either regular or occasional payments toward marketing costs or price promotions”*.³⁸ Since the magnitude of these payments is unknown to the public and kept confidential, the estimates that we provide are interesting on their own and are much more than just an input to the counterfactual analysis. In appendix A.2 we show that our empirical model shares the same setup to employ the identification strategy illustrated in Komarova, Sanches, Silva Jr., and Srisuma (forthcoming). We derive the closed-form identification of the adjustment costs in our model in appendix A.3. In addition, Komarova et al. (forthcoming) argue that the discount factor can generally be identified in a parametric model. In particular their Proposition 1 shows the discount factor is robust to normalisations in the adjustment costs. This implies that our estimates of the discount factor are independent of the restrictions we are imposing in the model in order to identify the vector of price adjustment costs.

Before moving to the description of our estimation strategy, a brief discussion on the identification of the CCPs is necessary. While our main source of variation is coming from the repeated play over time, to precisely estimate the CCPs we pool data from the four big supermarkets and include a supermarket fixed effect – see subsection 1.6.2 for arguments justifying this choice. In what follows we assume that the same equilibrium is played over time in all available markets conditional on the fixed effects.

³⁶In essence this result is a generalisation of Aguirregabiria and Suzuki (2014).

³⁷One can also prove identification under a different set of restrictions from the one we propose here, for instance it would suffice to assume that for every good the costs are symmetric (equal for changing prices from high to low and vice versa).

³⁸See The Guardian: <http://www.theguardian.com/business/2007/aug/25/supermarkets> (access on August 15, 2017).

1.6.2 Estimation

The estimation of the model relies on the following modelling two assumptions. First, we treat the big three manufacturers, Arla, Dairy Crest and Unilever, as the players involved in the dynamic pricing game. Each firm sets prices for the two different products they manufacture – see table 1.2. Given that the manufacturers in our model are market leaders with substantial bargaining power to influence price decisions, assuming that the manufacturers set prices is more realistic than assuming that supermarkets set prices and manufacturers are completely passive³⁹ – see Slade (1998). This assumption is also commonly employed in other empirical papers that use scanner data and do not have access to wholesale prices (e.g. in Nevo (2001b), Dubé et al. (2009)). Following this approach, marginal costs faced by the decision makers are assumed to include a retailer markup. Since marginal costs in our model are constant over time, this also means that retailer markups are fixed.

Second, the four retailing chains are treated as separate markets in which the same game is played independently. Store brand can be chosen by the consumers and is considered in the demand model, but as in Slade (1998), we believe that it is more appropriate to assume that supermarkets take the residual demand and do not act as active players.⁴⁰

There are also other reasons why we would not want to consider a game played between supermarket chains. First of all, there are by far more dimensions of competition between retailers than pricing of one category of goods. It would seem implausible to treat the profits of supermarkets in the butter/margarine category as separable from all their other activities (advertising, loyalty programs, opening of new outlets etc.). Even if we did assume so, the number of distinct products offered by each supermarket is much higher than the number of brands in the portfolios of each manufacturer. As each additional product significantly increases the computational complexity of the problem, we suspect that the solution to the problem would soon turn out to be infeasible.

Subsequently, the set of players is $\{Arla, Dairy\ Crest, Unilever\}$. They offer the following products: $\mathcal{J}_{Arla} = \{Anchor, Lurpak\}$, $\mathcal{J}_{Dairy\ Crest} = \{Clover, Country\ Life\}$, $\mathcal{J}_{Unilever} = \{Flora, ICBINB\}$. The actions available to Arla are $\mathcal{A}_{Arla} = \{(p_{Anchor}^H, p_{Lurpak}^H); (p_{Anchor}^H, p_{Lurpak}^L); (p_{Anchor}^L, p_{Lurpak}^H); (p_{Anchor}^L, p_{Lurpak}^L)\}$, and the sets of actions of the remaining players can be

³⁹A further argument is that in the presence of private labels that are known to yield higher margins for the retailers, supermarkets should have no incentives to price national brands aggressively (Meza and Sudhir, 2010). Lal (1990) argues that manufacturers use price promotions to limit store brand's encroachment into the market. Moreover, in the data we also observe smaller manufacturers, whose products are never on promotion. If we endowed supermarkets with all the bargaining power, it would be hard to justify why they decide to use different pricing strategies for products coming from different manufacturers. Finally, in two independent studies, Srinivasan et al. (2004) and Ailawadi et al. (2006) find that retailers hardly ever benefit from price promotions, and it is almost exclusively the manufacturers who can enjoy increased profits from temporary sales.

⁴⁰This means that the market share of the store brand is a payoff-relevant state variable for the remaining firms. For simplicity we assume that the price of own brand product does not change with time, otherwise it would be an additional dimension of the state space.

defined in a similar manner. To determine the actual values of p_*^H and p_*^L , we calculated the median weekly prices over the 200-week period, conditional on whether the product was on promotion or not (see table A.6.2 in the appendix).

Outline of the estimation procedures

The estimation procedure involves multiple stages which we outline below:

1. Use household-level data to estimate the demand system parameters (δ, γ, η) .
2. Plug $(\hat{\delta}, \hat{\gamma}, \hat{\eta})$ into (1.10) to obtain an estimate of $s_{jt}(a_{it}, \mathbf{a}_{-it}, \mathbf{s}_{t-1})$.
3. Use market-level data to estimate firms' conditional choice probabilities, i.e. obtain $\widehat{\Pr}_i(a_i = \ell | \mathbf{z})$ for all i and \mathbf{z} .
4. Use CCPs to get $\{\widehat{\mathbf{SC}}_i\}_{i=1}^N$ in closed form following Komarova et al. (forthcoming).
5. Plug the demand and SC estimates into the conditional value functions and estimate the discount factor, and parameters in the period profit function using the approach of Komarova et al. (forthcoming).

Two things are worth noting here. First, step 4 can be performed independently of step 5 since, the costs of adjusting prices, SC , are independent of β and period payoffs, π . Therefore the estimates of SC are robust to any potential misspecifications in π . Furthermore, we can significantly reduce the number of parameters in the model with minimal effort since these SC can be computed in closed form in terms of the CCPs. Secondly, step 5, differently from steps 1-4, depends on the estimation of expected value functions – the term $\int V_i(\mathbf{z}_{t+1}, \varepsilon_{t+1}) dQ(\varepsilon_{i,t+1})$ in equation (1.4). When the variables in the state space are continuous and/or the dimension of the state space is large, as is the case in our application, traditional methods used to compute value functions⁴¹ do not work appropriately. Consequently we employ a different approach and compute value functions using (flexible) parametric approximations – see Sweeting (2013), Fowlie et al. (2016) and Barwick and Pathak (2015). The algorithm used here is similar to the one used in Sweeting (2013) and is described in the appendix. In the remainder of this section, we discuss the estimation of the demand system, the estimation of the CCPs and finally the estimates of the dynamic parameters of our model.

⁴¹See Hotz et al. (1994), Aguirregabiria and Mira (2007), or Pesendorfer and Schmidt-Dengler (2008).

Demand

With a representative sample of \mathcal{H} households drawn from the population H and observed through T periods, and given our specification of consumer demand in (1.8) and the assumption that the unobserved choice shocks are independent over products, households and time, we can consistently estimate the parameters of the demand system using maximum likelihood. Since we are not modelling supermarket choice, the estimation samples consist of households that were recording butter/margarine purchases in only one of the supermarkets in the sample period. To check whether restricting the sample to *non-shoppers* does not induce non-random selection, we investigated the distribution of market shares in the full and restricted samples, to find no substantial differences apart from a moderately higher share of store-brand products at the expense of Arla's brands. Our demand estimates are shown in table 1.8.

Table 1.8: Demand estimates

	ASDA	MORRISONS	SAINSBURY'S	TESCO
δ_{Anchor}	-2.775 [-2.899; -2.651]	-2.883 [-3.043; -2.723]	-3.175 [-3.314; -3.036]	-3.836 [-3.910; -3.763]
δ_{Lurpak}	-2.064 [-2.193; -1.945]	-2.083 [-2.236; -1.930]	-2.862 [-3.004; -2.719]	-3.375 [-3.445; -3.306]
δ_{Clover}	-3.077 [-3.175; -2.980]	-2.757 [-2.860; -2.654]	-3.507 [-3.605; -3.409]	-3.814 [-3.866; -3.761]
$\delta_{Country\ Life}$	-2.930 [-3.051; -2.810]	-3.213 [-3.363; -3.063]	-3.792 [-3.934; -3.649]	-4.519 [-4.596; -4.442]
δ_{Flora}	-2.450 [-2.524; -2.375]	-2.334 [-2.415; -2.253]	-2.756 [-2.836; -2.676]	-3.075 [-3.117; -3.033]
δ_{ICBINB}	-2.516 [-2.580; -2.453]	-2.819 [-2.892; -2.745]	-3.369 [-3.447; -3.291]	-3.624 [-3.665; -3.583]
δ_{SB}	-2.903 [-2.970; -2.835]	-2.919 [-2.994; -2.845]	-2.772 [-2.844; -2.699]	-3.149 [-3.184; -3.115]
η	-0.745 [-0.799; -0.691]	-0.655 [-0.717; -0.594]	-0.356 [-0.414; -0.299]	-0.159 [-0.190; -0.128]
γ	3.037 [3.002; 3.071]	3.008 [2.967; 3.049]	2.931 [2.896; 2.967]	3.277 [3.256; 3.297]
N	104,946	71,294	102,939	280,828
pseudo- R^2	0.285	0.363	0.137	0.180

Note: Estimates obtained using the baseline definition of loyalty (only purchases in $t - 1$ matter). All parameters are significantly different from 0 at the 1% level. 95% confidence intervals reported below estimated coefficients, constructed using robust standard errors. *SB* denotes store brand.

In all four markets, consumer loyalty (measured by γ) appears to play a crucial role in determining consumers' choices. In fact, given the magnitude of the negative alternative-specific intercepts, within an acceptable range of prices, we can see that it is almost the loyalty effect alone making a purchase more attractive than the outside option. The price coefficients, η , are negative in all cases and reflect differences in the composition of each supermarket's clientele which are in line with common knowledge: people shopping in Asda and Morrisons are more price-sensitive

than Tesco and Sainsbury's customers.

We conducted two robustness checks: table A.7.1 shows results obtained under the assumption that a consumer can stay loyal for more than one period, while table A.7.2 presents what happens if we let γ vary across products. As expected, for the first case, loyalty parameters turn out to be lower than in the baseline case which we attribute to consumers' imperfect memory. When we let γ vary across products we find some differences – in particular consumers appear to be less loyal to margarine brands and, intuitively, to the supermarkets' own-label products. Even though we reject the null that all γ 's are equal, richer specification of the demand model does not lead to dramatic improvements in terms of the pseudo- R^2 , and therefore we treat the specification with just one loyalty coefficient as baseline. Furthermore, a homogenous loyalty parameter is more consistent with a psychological interpretation of consumer inertia, while with product-specific loyalty parameters one can think of them being endogenously affected by product characteristics or firms' marketing activities.

A natural question arising in demand estimation is whether we can obtain consistent estimates without controlling for potential endogeneity of prices. Due to the nature of the industry of interest, it is hard to imagine that there can be any product characteristics unobserved by the consumers and potentially correlated with prices that are not captured by product-specific intercepts. Moreover, due to the timing assumption in our model, we know that prices are set prior to the realisation of individual demand shocks. Hence, similarly to Griffith et al. (2017) and Pavlidis and Ellickson (2017), we can conclude that endogeneity of prices should not be a major issue.

Conditional choice probabilities

Prior to discussing structural estimation of the components of the payoff functions, we present reduced-form evidence in the form of multinomial logit estimates of conditional choice probabilities (table A.7.3 in appendix A.7). The covariates correspond to the components of \mathbf{z}_t in the theoretical model given our choice of players and products, as discussed in the previous section. Ideally, we should be estimating CCPs separately for each player in each of the 4 markets in our data using nonparametric methods. Even with a fully parametric specification, 200 periods quickly turn out not to be enough to precisely estimate 51 coefficients per player with enough precision. We therefore pooled data from four supermarkets and include fixed effects to account for the fact that equilibrium strategies might differ across markets.

We explore richer specifications that include higher order terms and interactions between state variables, and used post-LASSO as a regularisation tool to deal with the large number of param-

eters. While LASSO outperforms standard multinomial logit with unpooled data, the gains from sparsity were relatively modest when we used the entire sample, both in terms of in-sample fit and out-of-sample predictive performance. Moreover, squared market shares and interactions were usually not selected by the method. We also experimented with a random effects specification as well as discrete lagged market shares, finding no substantial differences.⁴²

The results in table A.7.3 are generally consistent with the descriptive evidence shown in section 1.4. In particular, we see that (i) players seem to take into account what they, rather than their competitors played last period and (ii) manufacturers' reactions to own past shares and competitors' past shares seem limited. Yet, the analysis of the coefficients in the multinomial model is much more involved. We refer the reader to section 1.4 for more details on these facts.

1.6.3 Dynamic parameters

Before showing the estimates of the dynamic parameters it would be helpful to recall the steps of our estimation procedure. First, in subsection 1.6.2 we show the estimates of the demand system. Second, based on the demand estimates the law that governs the transition of states was obtained. From the CCPs shown above and given the demand and state transition estimates the dynamic parameters of the model can be finally recovered. This last step will be the main object of this subsection.

■ **Adjustment costs.** We start the description of our results with the parameters capturing price adjustment costs. Following the identification arguments in subsection 1.6.1 we emphasise once again that price adjustment costs can be recovered directly from the CCPs, independently from the other model primitives. Our estimates are, therefore, robust to the specification of the demand system, state transitions, discount factor and the other parameters in period payoffs. Table 1.9 reports the estimated dynamic parameters reflecting the costs of switching from high to low prices.

All the estimates are negative and their relative magnitudes reflect differences in costs across products. Even though the figures for Arla have relatively large standard errors, they are quantitatively similar to the costs of other firms which are all statistically significant. Therefore, we believe that the large standard errors are an artifact of the sampling variation in the Arla data, rather than a feature of the industry making Arla different from the remaining players. There does not seem to be a lot of variation across supermarkets, which is consistent with the results reported in Slade (1998). Technically this result is reflecting the fact that the magnitude of supermarket fixed effects

⁴²All additional results described above are available upon request.

is relatively small in the CCPs.⁴³ All three Arla, Dairy Crest and Unilever, seem to be incurring very similar costs to change the prices for both of their products at the same time.

Table 1.9: Price adjustment costs.

	ASDA	MORRISONS	SAINSBURY'S	TESCO
Arla				
<i>SC_{Anchor}</i>	-2.177 (2.15)	-2.508 (2.56)	-2.511 (2.72)	-2.497 (2.33)
<i>SC_{Lurpak}</i>	-2.388 (2.05)	-2.438 (2.50)	-2.451 (2.6)	-2.441 (2.25)
<i>SC_{Both}</i>	-4.430 (3.00)	-4.746 (3.52)	-4.765 (3.60)	-4.745 (3.25)
DC				
<i>SC_{Clover}</i>	-2.589*** (0.68)	-2.584*** (0.78)	-2.583*** (0.88)	-2.582*** (0.83)
<i>SC_{Country Life}</i>	-2.149*** (0.64)	-2.154*** (0.79)	-2.131*** (0.9)	-2.155*** (0.85)
<i>SC_{Both}</i>	-4.536*** (0.84)	-4.544*** (0.95)	-4.547*** (1.06)	-4.557*** (1.01)
Unilever				
<i>SC_{Flora}</i>	-1.526** (0.50)	-1.612** (0.52)	-1.633** (0.51)	-1.627** (0.51)
<i>SC_{ICBINB}</i>	-2.251*** (0.63)	-2.445*** (0.61)	-2.446*** (0.6)	-2.451*** (0.6)
<i>SC_{Both}</i>	-4.111*** (1.67)	-4.291*** (1.64)	-4.311*** (1.58)	-4.319*** (1.52)

Note: Price adjustment costs are scaled by the variance of to the distribution of ϵ , which is assumed type-I extreme value with mean 0. Standard errors obtained using 100 bootstrap replications given in parentheses below the point estimates. Significance levels: *** 1%, ** 5%, * 10%.

The estimates in table 1.9 are scaled by the standard deviation of the payoff shocks and hence cannot be interpreted in monetary terms. To give to the reader a clearer picture of the magnitude of these costs we estimate the remaining parameters in the payoff function and compute the ratio between price adjustment costs and variable profits.

To provide this calculation, we need to estimate the remaining parameters in the variable profit function. Unfortunately, our attempts to estimate marginal costs and market size, H , produced implausible results.⁴⁴ Our explanation for this is the following: based on Sanches et al. (2016b) we can write the best response functions obtained from the solution of problem (1.4) as a linear function of the parameters in the period payoffs i.e., best response functions for player i can be represented as a system of the form $Y_i(\sigma, G, SC_i) = X_i(\sigma, G) \cdot \theta_i$, where $Y_i(\sigma, G, SC_i)$ is a

⁴³By excluding supermarket fixed effects from the CCPs the estimates of the other coefficients do not change significantly.

⁴⁴We experiment with various different methods to estimate these parameters. Instead of using parametric approximations of the value function we tried to estimate the parameters using forward simulations of the value function (Hotz et al., 1994). Alternatively we tried to discretise the state space compute value function using the closed form expression for the ex-ante value function (Pesendorfer and Schmidt-Dengler, 2008). All our attempts produced implausible estimates for marginal costs and H .

column vector⁴⁵ and $X_i(\sigma, G)$ is a matrix with 3 columns – same number of rows as $Y_i(\cdot)$. The vector $Y_i(\sigma, G, SC_i)$ depends on beliefs, transitions and on the estimates of adjustment costs obtained above; $X_i(\sigma, G)$ depends on beliefs and state transitions only. The vector θ_i contains the three parameters of interest, namely, marginal costs for both products of player i and H (scaled by the standard deviation of the payoff shock). In theory, this representation implies that the *sufficient* condition for the identification of these parameters is that $X_i(\sigma, G)$ has full column rank. However, we find the variables in $X_i(\sigma, G)$ were highly correlated with each other – correlations were above 0.95 – which can lead to very noisy estimates of the marginal costs and H . This is perhaps not very surprising given that we attempted to estimate a high-dimensional vector of cost parameters without instruments or any additional cost-side data. Therefore we decided to pursue a different strategy to estimate all the components of firms’ payoffs. Since marginal costs and variance of the shocks in our model are not *per se* dynamic parameters, we calibrate them and instead focus on estimating the discount factor which is more important for pricing dynamics. For marginal costs, we used the estimates obtained by Griffith et al. (2017) on a subsample of our data set. To select the best value H (scaled by the variance of the shocks), we estimate the discount factor for different possible values of H and rely on various measures of model fit to select the value that minimises the distance between observed and implied distributions of actions played by firms. We present our goodness-of-fit measures in section 1.7, and show the sensitivity of the remaining results to changes in H .

■ **Discount factor.** Since the results are very similar for all markets, for the sake of brevity, from this moment on we will only present results for the biggest (Tesco) and smallest (Morrisons) market in terms of annual sales. We present the results for different values of H – using the calibrated marginal costs. All values of H outside the range of values shown in the table provided much worse fit to the data and are therefore omitted. The discount factors estimated using the method outlined in appendix B3 and section 3 of Komarova et al. (forthcoming) are presented in table 1.10. Our results show that firms are forward-looking, with the discount factors close to the typical values assumed in models calibrated using weekly data. We conclude the forward-looking behaviour of manufacturers is a crucial component of pricing models.

■ **Relative magnitudes of adjustment costs.** In table 1.11, we relate the present value of adjustment costs for all firms to their variable profits. The results are presented for different assumptions about H . As seen in table 1.11, the magnitudes of promotional costs are relatively large. Over

⁴⁵Number of rows is equal to the number of states times the number of possible actions minus one.

Table 1.10: Estimated discount factors.

H	MORRISONS β	TESCO β
0.50	0.9807*** (0.04)	0.9970*** (0.01)
1.00	0.9815*** (0.03)	0.9970*** (0.01)
2.00	0.9811*** (0.02)	0.9958*** (0.01)
3.00	0.9790*** (0.02)	0.9936*** (0.01)
4.00	0.9757*** (0.01)	0.9914*** (0.01)
5.00	0.9708*** (0.01)	0.9895*** (0.01)
6.00	0.9620*** (0.01)	0.9878*** (0.01)
7.00	0.9472*** (0.02)	0.9860*** (0.01)
8.00	0.9299*** (0.02)	0.9838*** (0.01)
9.00	0.9079*** (0.03)	0.9805*** (0.01)

Note: Results shown for different values of market size scaled by the variance of the shock, under the assumption that this value is the same for all firms, but potentially different across markets. Standard errors obtained using 100 bootstrap replications provided in parentheses below the point estimates. Significance levels: *** 1%, ** 5%, * 10%.

the horizon of 200 weeks, firms have to sacrifice approximately 24-34% of their variable profits in order to be able to charge promotional prices in some periods. These estimates are in line with existing evidence in the macro literature.⁴⁶ Since in absolute terms, these costs were very similar across players and the firms we considered are the market leaders, one can imagine that these costs constitute a much bigger fraction of the profits of smaller companies and local dairies, effectively restricting the scope of their promotional activities. This is consistent with what we observe in the data on smaller brands which were considerably less frequently on promotion.⁴⁷ Our analysis therefore shows price adjustment costs may have important implications for market structure.

In summary:

1. Table 1.8 suggests that loyalty is an important driver of consumer decisions, as previously pointed out by our descriptive evidence (tables 1.3 and 1.4 in section 1.4);
2. Our adjustment cost estimates represent a large fraction of manufacturers' payoffs. This

⁴⁶For instance, Levy et al. (1997) use store-level data to study the process of changing prices. They find that these costs represent 35.2% of net margins of retailers. Using the same approach Dutta et al. (1999) study price adjustment costs of a large US drugstore chain. Findings are similar to the findings of Levy et al. (1997). Price adjustment costs – physical and labor costs of changing prices – amounts to 27.08% of net profit margins. In addition to physical costs involved in price adjustment processes Zbaracki et al. (2004) quantify managerial and customer costs of price adjustment using data from a large industrial manufacturer. Managerial costs are defined as the managerial time and effort spent with pricing decisions; customer costs are defined as the costs of communicating new prices to consumers. Price adjustment costs adds up to 20.03% of company's net margins. It is worthwhile mentioning that all these evidence are direct, in the sense that they were obtained directly from accounting data.

⁴⁷See figure A.6.1 in appendix C.

Table 1.11: Magnitude of adjustment costs.

H	MORRISONS			TESCO		
	Arla	DC	Uni	Arla	DC	Uni
0.5	-34.70%	-32.98%	-30.81%	-33.49%	-33.46%	-30.36%
1.0	-34.37%	-32.71%	-30.23%	-33.16%	-33.26%	-29.89%
2.0	-33.78%	-32.22%	-29.22%	-32.52%	-32.93%	-28.97%
3.0	-33.29%	-31.74%	-28.30%	-31.84%	-32.61%	-28.11%
4.0	-32.79%	-31.26%	-27.48%	-31.11%	-32.27%	-27.30%
5.0	-32.35%	-30.72%	-26.73%	-30.44%	-31.95%	-26.50%
6.0	-32.09%	-30.17%	-26.07%	-29.72%	-31.62%	-25.73%
7.0	-31.99%	-29.41%	-25.50%	-29.04%	-31.29%	-25.01%
8.0	-31.80%	-28.61%	-24.94%	-28.30%	-30.96%	-24.27%
9.0	-31.61%	-27.60%	-24.33%	-27.52%	-30.66%	-23.58%

Note: The numbers in the table are ratios of adjustment costs to variable profits for each firm in two different supermarkets. Both components of the payoff are calculated as average present values for 200 periods, averaged across 1000 simulated paths.

can help explain the price rigidities found in the descriptive evidence in section 1.4. The magnitudes of our adjustment costs are in line with (direct) evidence found in the macro literature and, judging by their relative importance on manufacturers' payoffs, it is likely that price adjustment costs have implications for the structure of this market.

3. The estimates of the discount factor point to a high degree of forward-looking behaviour and their size is similar to the β 's typically assumed in the literature.

Next we examine the goodness-of-fit of the model and propose a series of counterfactual studies. The key purpose of these counterfactuals is to study the implications of consumer loyalty on price dynamics when price adjustment is costly.

1.7 Model fit and counterfactuals

This section begins with us justifying our choice of the grid used to calibrate H , as well as present some measures of model fit. We will turn to our two counterfactual studies. In light of the findings in sections 1.4 and 1.6, it seems that both, consumer loyalty and price adjustment costs, are fundamental to the understanding of the price process in this market. We want now to quantify the implications of consumer loyalty on price dynamics in the presence of price adjustment costs. Intuitively, price adjustment costs may mitigate the incentives of firms to invest in a broader consumer base through price promotions. To benefit from this of type strategy firms have to bear not only temporary profit losses steaming from temporary price reductions but also the price adjustment cost itself. Our counterfactuals will serve to illustrate this intuition. Second, the estimates in subsection 3.5 indicate that price adjustment costs are substantial. An additional objective of this section is to quantify the importance of price adjustment costs to consumers and firms. This

last exercise may help us to understand how investments in practices and technologies that aim to reduce price adjustment costs affect prices, profits and consumer surplus.

1.7.1 Model fit

This subsection analyses the fit of the model and describes the arguments that guided our choices of H . We select H 's for each market by examining two measures of model fit (see table 1.12). To calculate these measures, we take the vector of market shares observed in the first period of our data as initial conditions, and simulate the model 199 periods ahead using the equilibrium CCPs. We repeat the simulation 1,000 times and compare simulated and real data to calculate: (i) the sum of absolute differences between the fractions of periods in which each action was played by the three firms; (ii) sum of absolute differences between market shares of all brands.

While the numbers in the table may not have an obvious interpretation, it is clear that we want to minimise both of them. For both markets, values of H higher than 9 yielded much worse fit. Moreover, the expected payoffs quickly reach (computer) infinity as H increases making the computation of counterfactual equilibrium infeasible. For the values of $H \in \{0.5, \dots, 9\}$, we observe that in general, lower values give rise to a better fit of the market shares, though the differences are very small. We observe more noticeable differences for the fit of actions, and hence rely on this metric for our choice of the best model ($H = 8$ for Morrisons and $H = 3$ or $H = 4$ for Tesco). In principle we could also refine the grid around these values, but that would only affect the computational time, without having any serious qualitative impact on our remaining results.

Table 1.12: Measures of model fit.

H	MORRISONS		TESCO	
	Actions	Shares	Actions	Shares
0.5	0.803	0.021	0.982	0.011
1.0	0.802	0.022	0.984	0.011
2.0	0.790	0.022	0.984	0.012
3.0	0.775	0.022	0.980	0.012
4.0	0.750	0.022	0.981	0.012
5.0	0.716	0.023	0.990	0.012
6.0	0.673	0.023	1.007	0.012
7.0	0.617	0.024	1.038	0.013
8.0	0.591	0.024	1.079	0.013
9.0	0.686	0.025	1.131	0.013

Note: For both supermarkets, two measures of model fit are reported for different calibrations of H . The first one (second and fourth column) is the sum of absolute differences between the fractions of periods with a given action being played observed in the data and simulated from the equilibrium of the model. The second statistic, reported in columns 3 and 5, measures the absolute difference between observed and simulated market shares. Data from the equilibrium of the model were simulated 1,000 times, 199 periods ahead, using the state observed in week 1 of the data as initial conditions.

For the models providing best fit, we decompose the above measures of fit by firm and brand,

respectively (see figures A.7.1 and A.7.2 in appendix D). The model does a good job fitting market shares and predicting firms' pricing behaviour. Only for Arla, we underestimate the number of periods in which one of the brands is on sale. For the other firms we manage to replicate the distribution of actions quite accurately.

1.7.2 Counterfactuals

Equipped with the estimates of the payoffs, we can now answer the questions posed at the beginning of the chapter. Namely, we seek to understand (i) how consumer loyalty affects price dynamics in the presence of price adjustment costs and (ii) how price adjustment costs affect firms' profits, equilibrium prices and consumer welfare.

■ **Consumer loyalty under price adjustment costs.** To examine the first question we simulate the pricing game using the estimates of price adjustment costs (see table 1.9), our calibrated values for H (according to table 1.12), the corresponding discount factor estimates (see table 1.10) and different values for γ , which is the parameter that captures consumer loyalty in our model. We redo the same exercise setting $SC_i = 0$ for all firms and compare equilibrium prices (averaged across the 6 brands) produced by the models with and without price adjustment costs. To compute equilibrium prices we solved the model using the value function approximation method described in the technical appendix. Starting from the state vector observed at the first week in our sample we simulate the model 199 periods ahead 1000 times and compute average prices across periods and simulations. We solve the model using different initial guesses for the vector of equilibrium probabilities, to detect possible multiplicity of equilibria, finding our algorithm to converge to the same equilibrium every time.

Table 1.13 shows the results for Morrisons and Tesco. The first column has the factor that we use to scale the parameter capturing consumer loyalty (γ in table 1.8). Columns 2 and 3 show average prices and the percentage difference of prices between the model in the corresponding row and the model without consumer loyalty (i.e. the model in the first row) for the MPE simulations where price adjustment costs are set to zero. The two subsequent columns have the same statistics for the models with price adjustment costs. The last column has the price variation between the models with and without price adjustment costs.

The table shows some interesting results. First, increases in consumer loyalty are associated with increases in equilibrium prices in the models with and without price adjustment costs. This observation holds for both supermarkets. For lower levels of consumer loyalty the effects of increases in the loyalty factor on prices are relatively small (but still positive). When the levels

Table 1.13: Implications of consumer loyalty with and without price adjustment costs.

Loyalty factor	SC= 0		Estimated SC		Price SC Price SC=0
	Price	Difference	Price	Difference	
MORRISONS					
0.00	1.750	-	1.797	-	2.69%
0.25	1.750	0.01%	1.798	0.02%	2.70%
0.50	1.751	0.03%	1.799	0.07%	2.73%
0.75	1.751	0.07%	1.800	0.18%	2.80%
1.00	1.753	0.16%	1.805	0.41%	2.94%
2.00	1.811	3.49%	1.896	5.47%	4.66%
3.00	1.858	6.16%	1.944	8.17%	4.63%
TESCO					
0.00	1.740	-	1.754	-	0.80%
0.25	1.741	0.00%	1.755	0.01%	0.80%
0.50	1.741	0.01%	1.755	0.04%	0.82%
0.75	1.741	0.03%	1.756	0.10%	0.87%
1.00	1.742	0.07%	1.758	0.22%	0.95%
2.00	1.760	1.14%	1.816	3.50%	3.15%
3.00	1.769	1.62%	1.853	5.63%	4.77%

Note: Columns labeled “Price” contain average prices (across the 6 branded products); columns labeled “Difference” contain the percentage difference between prices in the corresponding row with respect to the price obtained from the model where the loyalty factor is zero (i.e. prices in the first row); the last column has the price difference between the models with and without price adjustment costs in the corresponding row. The figures were obtained by simulating the two models according to MPE choice probabilities 200 periods ahead, and averaging across 1,000 simulation paths.

of consumer loyalty are already high, increases in the loyalty factor lead to a increase in prices. These patterns are similar to those found in Dubé et al. (2009) with one important exception. In Dubé et al. (2009) prices initially fall for lower consumer loyalty levels, whereas in our case firms seem to have an insufficient incentive to invest in building up their consumer base.

Second, the consequences of consumer loyalty for prices are more pronounced in the model with price adjustment costs. For example, in Tesco, a change in the loyalty factor from 0 to 3 causes a price variation of 1.62% in the model where price adjustment costs are zero and of 5.63% in the model with price adjustment costs. The same conclusion holds for Morrisons and for each brand separately. The differences in the magnitudes of these effects between Tesco and Morrisons may be explained by differences in H . In particular this parameter is much smaller for Tesco than Morrisons’, which suggests that changes in consumer switching costs will have more relevant implications in Morrisons than in Tesco. Our conclusion is that price adjustment costs may act as an additional barrier for firms that want to invest in consumer loyalty through temporary price reductions.

Finally, in line with the descriptive regressions in section 1.4, price adjustment costs appear to be more important to explain price dynamics than consumer loyalty. From our baseline estimates (rows in bold) the inclusion of price adjustment costs in the model leads to a increase of 3% in

average prices for Morrisons and of 1% for Tesco. This contrasts with the effects of consumer loyalty on prices. In the model with price adjustment costs, an increase in the loyalty factor from zero (no consumer loyalty) to one (baseline estimates of consumer loyalty) causes a price increase of approximately 0.4% in Morrisons and of 0.2% in Tesco.

■ **Price adjustment costs, profits and consumer surplus.** The results in table 1.13 also suggest that, despite their magnitudes, price adjustment costs have a small effect on final prices. Next we provide further evidence on this result. We start with an analysis of the effects of price adjustment costs on profits and consumer surplus. The results of this study are shown in table 1.14.

To construct this table we compute the percentage differences between baseline (model with price adjustment costs) and counterfactual (model without price adjustment costs) profits and market shares of each manufacturer and consumer surplus. While we focus on the results for the two calibrations of H which provide best fit of the model, the table includes also welfare measures for alternative values of the parameter to show that our main qualitative conclusion is robust to the choice of H . Not surprisingly, eliminating this type of friction has a large positive effect for firms' profits, ranging from 50 to almost 75%. This is considerably more than the magnitude of the promotional costs alone (see table 1.11), which represent 20-30% of firms variable profits. This difference is mainly explained by an increase in the expected value of the profitability shock for the firms, i.e. the (conditional) expectation of term $\sum_{\ell \in \mathcal{A}_i} \varepsilon_{it}(\ell) \cdot \mathbf{1}(a_{it} = \ell)$ in equation (1.2). Clearly, an important implication of this finding is that investments in managerial practices and technologies that reduce price adjustment costs generate large returns for firms – see, for example, Basker (2012) and Ellison et al. (2015) for other studies on the effects of process innovation of this type on profits. Also, as alluded in section 3.5, these findings suggest that price adjustment costs may have, in the long-run, considerable influence on market structure. Without price adjustment costs potential entrants will expect considerably higher profits in the long-run. This effect might, in the end, induce the entry of new competitors in the industry.

Consumer surplus, on the other hand, increases only by a modest percentage when price adjustment costs are removed from the model. Competition in this market appears to be limited, which means that incumbents do not have incentives to pass the cost reduction to consumers. To understand this result better, we further decompose our findings and look at other margins in table 1.15.

In both supermarkets, under costless price adjustment, we observe an increase in the number of weeks where each firm has at least one of its brands on promotion. However, the drop in the average long-run price paid by the consumers ranges only between 1 and 6p, which explains the

Table 1.14: Counterfactual results with $SC = 0$.

H		MORRISONS			TESCO		
		Arla	DC	Uni	Arla	DC	Uni
0.5	Δs	0.63%	0.59%	0.30%	0.08%	0.06%	0.04%
	$\Delta \Pi$	82.87%	75.38%	64.63%	78.88%	76.55%	63.41%
	ΔCS		0.44%			0.05%	
2.0	Δs	0.88%	0.72%	0.61%	0.30%	0.13%	0.14%
	$\Delta \Pi$	79.49%	72.99%	60.22%	75.64%	74.82%	72.89%
	ΔCS		0.70%			0.18%	
4.0	Δs	1.47%	1.14%	1.20%	0.63%	0.26%	0.27%
	$\Delta \Pi$	76.24%	70.01%	55.68%	71.49%	72.89%	55.08%
	ΔCS		1.25%			0.37%	
6.0	Δs	2.41%	2.14%	1.84%	0.95%	0.38%	0.38%
	$\Delta \Pi$	74.34%	67.16%	52.37%	67.55%	70.96%	51.14%
	ΔCS		2.04%			0.54%	
8.0	Δs	3.97%	3.80%	2.77%	1.25%	0.51%	0.48%
	$\Delta \Pi$	74.51%	64.16%	50.52%	63.80%	69.14%	47.67%
	ΔCS		3.27%			0.71%	

Note: Numbers in the table are percentage differences between the counterfactual scenario and the baseline model in: average market share (Δs), firm profits ($\Delta \Pi$) and consumer surplus (ΔCS). The figures were obtained by simulating the two models according to MPE choice probabilities 200 periods ahead, and averaging across 1,000 simulation paths.

aforementioned modest increase in consumer surplus. The most important difference between the baseline scenario and the counterfactual is in the duration of promotional periods – the lack of adjustment costs makes firms choose shorter, albeit more frequent, periods of temporarily reduced prices. We would therefore no longer be observing the persistence of prices which we spotted in the original data, though this difference turns out to have very little effect on consumer surplus.

The results of this counterfactual can also be interpreted as partial equilibrium response to a ban on promotional fees. While such regulation has not been proposed in the UK yet, similar policies have been implemented in some countries to increase the degree of transparency in the retailer-manufacturer relationships.⁴⁸ Our results indicate that such a regulation would have a modest impact on consumer welfare and would simply shift the profits from retailers to manufacturers in the vertical channel. This part of the result should be interpreted with caution because we are not analysing the general equilibrium response of the downstream firms (supermarkets).

⁴⁸See The Economist: <http://www.economist.com/news/business/21654601-supplier-rebates-are-heart-some-supermarket-chains-woes-buying-up-shelves> (accessed on August 15, 2017): "Some countries have tried to protect consumers by making rebates illegal. Poland banned them in 1993 (...). And in 1995 America banned them on alcoholic drinks (...). However, progress towards eliminating them on all products in America stalled after the Federal Trade Commission (FTC) concluded in 2001 that more research on them was needed before it could take any further action". (access March 8, 2018).

Table 1.15: Decomposition of main counterfactual results.

		MORRISONS: $H = 8$		TESCO: $H = 4$	
		Baseline	Counterfactual	Baseline	Counterfactual
Arla	No promotions				
	◊ <i>Frequency</i>	37.82%	26.50%	31.39%	26.56%
	◊ <i>Avg. duration</i>	3.08	1.36	2.79	1.36
	One promotion				
	◊ <i>Frequency</i>	46.88%	49.91%	49.09%	49.97%
	◊ <i>Avg. duration</i>	2.43	1.33	2.49	1.33
	Two promotions				
	◊ <i>Frequency</i>	15.29%	23.59%	19.51%	23.47%
	◊ <i>Avg. duration</i>	2.07	1.31	2.24	1.31
	\bar{p} Anchor	£2.25	£2.20	£2.23	£2.21
\bar{p} Lurpak	£2.45	£2.39	£2.38	£2.34	
Dairy Crest	No promotions				
	◊ <i>Frequency</i>	35.96%	26.19%	28.81%	25.70%
	◊ <i>Avg. duration</i>	2.88	1.36	2.40	1.33
	One promotion				
	◊ <i>Frequency</i>	47.80%	49.97%	49.14%	49.90%
	◊ <i>Avg. duration</i>	2.37	1.33	2.40	1.33
	Two promotions				
	◊ <i>Frequency</i>	16.23%	23.83%	22.05%	24.40%
	◊ <i>Avg. duration</i>	2.05	1.32	2.30	1.33
	\bar{p} Clover	£1.49	£1.43	£1.48	£1.46
\bar{p} Country Life	£2.14	£2.10	£2.09	£2.08	
Unilever	No promotions				
	◊ <i>Frequency</i>	38.46%	27.99%	30.14%	26.62%
	◊ <i>Avg. duration</i>	2.71	1.39	2.37	1.37
	One promotion				
	◊ <i>Frequency</i>	47.72%	50.26%	49.95%	50.04%
	◊ <i>Avg. duration</i>	2.13	1.34	2.17	1.33
	Two promotions				
	◊ <i>Frequency</i>	13.83%	21.75%	19.92%	23.33%
	◊ <i>Avg. duration</i>	1.77	1.29	2.00	1.31
	\bar{p} Flora	£1.25	£1.21	£1.28	£1.26
\bar{p} ICBINB	£1.05	£1.01	£1.07	£1.06	

Note: The table compares various summary statistics in the baseline scenario where price adjustment is costly and in the counterfactual with no promotional costs. For each firm, we present simulated frequency and duration of different actions (first six rows), and average long-run prices of each brand, weighted by market shares, denoted as \bar{p}_* .

1.8 Summary and conclusions

This chapter analysed multiproduct pricing in an environment where consumers exhibit inertia in their choices and oligopolists might be facing costly price adjustments. Based on the empirical observation that the distribution of retail prices has only a few mass points, we cast the problem as a dynamic discrete game and analysed pure strategy Markov perfect equilibria. We employ recent identification results by Komarova et al. (forthcoming) to arrive at a tractable estimation strategy, which allows us to estimate the cost of adjusting prices and firms' discount factor.

We apply the model to the UK butter and margarine industry and estimate the structural parameters using a detailed scanner data set. First, our estimates of price adjustment costs show that

firms pay between 24-34% of their variable profits to change their prices. The magnitudes of these costs are in line with (direct) evidence found in the macro literature – see, for instance, Levy et al. (1997), Dutta et al. (1999), Zbaracki et al. (2004). Second, using the methodology proposed in Komarova et al. (forthcoming) we also estimated the discount factor of butter and margarine producers. Our discount factor estimates are within the range of values commonly assumed in other dynamic pricing studies – see Dubé et al. (2009) and Pavlidis and Ellickson (2017). This result implies that firms’ forward-looking behaviour is a critical component of pricing models.

We use the model to understand the effects of consumer loyalty on prices when price adjustment is costly. Our first counterfactual exercise finds that price adjustment costs dampen firms incentives to invest in consumer loyalty, which exacerbates potentially negative effects of consumer switching costs on prices. Our second counterfactual study shows that price adjustment costs also have important effects on firms profits. By removing price adjustment costs from the market we observe a significant increase in profits but little effect on prices and consumer welfare. Given their magnitudes, it is very likely that price adjustment costs may have consequences for market structure. Smaller firms may not have the capacity to pay these costs to lower their prices frequently which, in turn, lowers their ability to enter and compete in this market. A systematic investigation of price adjustment costs on entry and exit dynamics seems to be an interesting topic for future research.

Chapter 2

Identification and Estimation of a Search Model: A Procurement Auction Approach

2.1 Introduction

Many theoretical models have been developed to explain price dispersion of homogeneous products relying on the notion that search is costly for consumers; see the survey of Baye, Morgan, and Scholten (2006). The econometrics literature has exploited the structure imposed by search models to show nonparametric identification, relying on the *fixed sample*¹ framework studied in Burdett and Judd (1983). In this model, *ex ante* identical firms compete by setting prices in a complete information environment, while a continuum of consumers choose how many prices to sample. Price dispersion arises in equilibrium as firms employ a mixed strategy Nash pricing rule in equilibrium. Hong and Shum (2006) alter the original model by assuming that consumers differ in their search cost which is drawn from a continuous distribution. They then use the indifference condition which defines the mixed strategy NE to show that this distribution can be identified from data on prices alone. When the dataset available is limited to a single market, only a finite number of points of the search cost distribution can be identified. The identification of the search cost distribution over its entire support is possible, for instance, when we have more data on prices from different markets, as shown by Moraga-González, Sándor, and Wildenbeest (2013).

Since the main goal of Burdett and Judd (1983) was to establish the existence of price dis-

¹In a fixed sample (nonsequential) search consumers decide simultaneously how many firms to search from. This stands in contrast to sequential search. Some recent empirical studies found that nonsequential search models provide a better approximation to consumers' search behaviour observed in real life (De los Santos, Hortaçsu, and Wildenbeest (2012), Honka and Chintagunta (2017)).

persion in the simplest possible environment with no heterogeneity on either side of the market, the model might not be very well suited to most empirical applications. In this chapter we therefore consider a more general setting where firms can have heterogeneous costs of production. Our model builds on the theoretical framework proposed by MacMinn (1980), where firms have independent, private marginal costs.² This approach leads to a game of incomplete information played between firms that resembles a procurement auction. The subsequent solution concept is a Bayesian-Nash pure strategy equilibrium, in which firms adopt monotone pricing rules. Therefore, price dispersion observed in our model can be attributed to both search and *ex ante* differences in firms' productivity. We aim to provide a general framework for a structural econometric analysis of such a search model.

The key contribution of this chapter is to provide a theoretical, both in terms of economics and econometrics, treatment for analysing an empirical search model and lay out a corresponding estimation methodology. We characterise the equilibrium of the model, provide conditions under which one can identify the distributions of consumers' search costs and firms' marginal costs. Finally, we propose nonparametric estimators for all of the identified objects in the model and provide asymptotic properties of these estimators when appropriate data are available.

Our identification strategy differs from those employed to study models in the spirit of Burdett and Judd (1983). The insight of Hong and Shum (2006) uses the constancy condition imposed by a mixed strategy equilibrium to identify the distribution of consumers' search cost. We do not have such a restriction to exploit with the pure strategy BNE solution concept. Therefore, in addition to price, we require data on a variable other than price to identify the proportions of consumer search. We then assume the regression of this variable on price to be related to the proportions of consumer search through a particular semiparametric index restriction. These proportions need to be recovered in the first place, as they appear in the firms' optimal pricing rules, which are subsequently used to identify the distribution of marginal costs. Following our identification steps, we propose a companion two-step estimation procedure:

Step 1. The proportions of consumer search are estimated. When the index specification is linear, our estimator admits an OLS-type, closed-form solution.

Step 2. The firms' marginal costs are estimated. These generated variables are then used to construct a nonparametric estimator for the probability density function of the marginal costs in a similar fashion to Guerre, Perrigne, and Vuong (2000,

²Similarly to us, Benabou (1993) considers a search model with bilateral heterogeneity. He does not, however, present any results on nonsequential search, so we are more likely to see our results as an extension of MacMinn (1980), instead of a fixed sample version of his model.

hereafter GPV).

Despite its seemingly natural scope for applications³, we are not aware of any theoretical work that considers our search model previously. We directly extend MacMinn's partial equilibrium analysis⁴ of fixed sample search model to a full equilibrium one with an arbitrary continuous distribution of marginal costs defined over a compact support.⁵ We build on his insight that makes the connection between the search and procurement auction models. The pricing problem of each firm can be seen as a first price procurement auction problem with random participation; the number and identity of bidders are stochastic. We characterise an equilibrium that generates a continuous price distribution.

The model of search that is closest to the one we consider in this paper can be found in a recent empirical study by Salz (2017). A version of our model can in fact be seen as a special case of his. However, as an econometric problem, our search problems are not nested. In particular his identification strategy is not applicable to our model.

Salz studies the trade-waste market in New York City. In his model buyers (consumers) can haggle (search) directly with carters (firms), or use a broker who has access to a group of carters. The haggling part is the same as our search problem. A broker acts as a clearinghouse where a standard procurement auction game with known number and identity of bidders is played. Salz assumes an equilibrium exists in his model. Importantly, Salz's identification strategy relies on the assumption that a broker *always* exists; see his Assumption 1. He also assumes both carters that can be searched and those who participate with brokers have the same cost distribution.⁶ Therefore he can identify the firm's cost distribution using the procurement auction data from the brokers independently of the search mechanism. The identification for the remaining components of his model subsequently relies on this.

Brokers do not exist in our model. We emphasize that we are not being critical of Salz's approach. His model captures important features of many real world industries. Nevertheless brokers, or other clearinghouse facilities, need not exist in many other markets. For these pure search models we show identification is possible with additional data. Our key identifying assumption

³Our model generates price dispersion in a transparent manner through heterogenous marginal costs. A mixed strategy solution is harder to interpret.

⁴MacMinn derives firms' best response functions taking search behaviour as given and proves that an equilibrium distribution of prices exists for an arbitrary number of prices sampled. The paper does not, however, directly consider simultaneous determination of search behaviour and optimal pricing. Hence, in line with other authors (Pereira, 2005), we refer to his paper as partial equilibrium analysis.

⁵As opposed to the results in MacMinn (1980), which are derived assuming that firms' costs are uniformly distributed.

⁶Salz assumes there are two types of carters. H(igh) and L(ow) cost types. Both types are present in both the broker and search markets. A carter that participates in both markets generally will bid differently during the auction and haggling process.

involving the semiparametric index restriction is empirically motivated. It contains as a special case the assumption that some observable market outcomes are proportional to the probabilities of firms completing a sale as consumers search in expectation. Natural candidates for such variable could be market shares or sales figures. This idea is identical to linking market shares to the choice probabilities, which is the starting point for the identification argument used in the study of differentiated products markets in the IO literature (see Berry and Haile (2014)).

In terms of econometrics, the estimation of the demand side of our model is relatively straightforward. The estimators for the demand parameters are smooth functionals of the empirical process of observed prices and will converge at a parametric rate; cf. Sanches, Silva Junior, and Srisuma (2016a). The estimation of the distribution of marginal cost is more challenging. Following the tradition set by GPV for nonparametric analysis of auction models, we focus on density estimation and study its uniform convergence rate. The marginal cost PDF turns out to be the most difficult object to estimate in our model.

We employ the same estimation strategy as GPV. We first use the observed prices to generate the latent, or pseudo-, marginal costs and then perform nonparametric estimation using the generated variables. To this end we establish some key relations between the density function of the observed and latent variables in our model. These findings are not just for theoretical interest but have important practical implications. The crucial one is that the density of the observed price generally asymptotes to infinity as the price approaches the upper bound of its support. Estimating a density function with a pole requires particular care as standard kernel estimation techniques are only suitable when the underlying density is assumed to be bounded on its support. For this we characterise the behaviour of the price density at the upper boundary and suggest a transformation method that eliminates the boundary issue (cf. Marron and Ruppert (1994)). However, a slower uniform convergence rate in the neighbourhood of the pole than other part of the support is a necessary feature. We show our estimator has the same convergence rate as the GPV estimator on any compact inner subset of the support. The density over an appropriately expanding support will converge at a slower rate depending on the speed of the support expansion. We can make the convergence rate to be arbitrarily close to the optimal convergence rate derived in GPV's auction problem.

The rest of this chapter proceeds as follows. Section 2 presents the model and characterises the equilibrium of the game. Section 3 presents our constructive identification strategy. Section 4 contains the theoretical results. Section 5 discusses ideas for extensions. Section 6 presents a simulation study.

2.2 Model

We consider a model in which there is a unit mass of consumers and a finite number of firms. Each consumer has an inelastic demand for a single unit of a good supplied by the firms. Consumers differ by search costs and employ a nonsequential search strategy to find the lowest price, at which they purchase the product. We next formally introduce the elements of the game.

2.2.1 Supply Side

There are I firms. Let $\mathcal{I} \equiv \{1, \dots, I\}$. Firm i draws a marginal cost of production R_i . R_i is assumed to be a continuous random variable supported on $[\underline{R}, \overline{R}] \subset \mathbb{R}$. We denote its cumulative distribution function (CDF) by $H(\cdot)$. The marginal costs of firms are independent from each other. Firm i then faces the following decision problem:

$$\begin{aligned} & \max_p \Lambda(p, R_i; \mathbf{q}), \text{ where} \\ \Lambda(p, R_i; \mathbf{q}) &= (p - R_i) \sum_{k=1}^I q_k \frac{k}{I} \mathbb{P}[P_{(1:k-1)} > p]. \end{aligned}$$

Here $\mathbf{q} = (q_1, \dots, q_I)^\top$ denotes a vector containing $(q_k)_{k=1}^I$ where q_k denotes the proportion of consumers searching k firms. For a given k , $\frac{k}{I}$ is the number of combinations that firm i gets included when k firms are sampled.⁷ We use $P_{(k:k')}$ to denote the k -th order statistic from k' i.i.d. random variables of prices with some arbitrary distribution; $P_{(1:k-1)}$ denotes the minimum of such $k-1$ prices. Here we implicitly assume that all firms have equal probability of being found and thus the game is symmetric. We discuss how this assumption can be relaxed in Section 5.2.

Firm's Best Response

We assume there exists a candidate for an optimal symmetric pricing strategy $\beta : [\underline{R}, \overline{R}] \rightarrow [\underline{P}, \overline{P}] \subset \mathbb{R}$ with the following properties: (i) β is strictly increasing; (ii) $\beta(\overline{R}) = \overline{R}$, which is the free entry condition imposing that $\overline{P} = \overline{R}$.

Let \mathbb{S}^{I-1} denote a unit simplex in \mathbb{R}^{I+} . For any $\mathbf{q} \in \mathbb{S}^{I-1}$, we can define $\Lambda^*(\cdot; \mathbf{q})$ to be the value function for a representative firm when all players are assumed to employ a strictly

⁷Let $C_k^I \equiv \frac{I!}{(I-k)!k!}$ denote the combinatorial number from choosing k objects from a set of I . Then $C_{k-1}^{I-1}/C_k^I = \frac{k}{I}$.

increasing optimal pricing strategy that we denote by $\beta(\cdot; \mathbf{q})$. We denote $\beta^{-1}(\cdot; \mathbf{q})$ by $\xi(\cdot; \mathbf{q})$.

$$\Lambda^*(r; \mathbf{q}) = (\beta(r; \mathbf{q}) - r) \sum_{k=1}^I q_k \frac{k}{I} (1 - H(\xi(\beta(r; \mathbf{q}); \mathbf{q})))^{k-1}.$$

Then by the envelope theorem (Milgrom and Segal (2002)),

$$\begin{aligned} \left. \frac{d}{dr} \Lambda^*(r; \mathbf{q}) \right|_{r=R} &= - \sum_{k=1}^I q_k \frac{k}{I} (1 - H(R))^{k-1}, \text{ and} \\ \Lambda^*(\bar{R}; \mathbf{q}) - \Lambda^*(R; \mathbf{q}) &= - \sum_{k=1}^I q_k \frac{k}{I} \int_{s=R}^{\bar{R}} (1 - H(s))^{k-1} ds. \end{aligned}$$

Thus for any r ,

$$\beta(r; \mathbf{q}) = r + \frac{\sum_{k=1}^I q_k k \int_{s=r}^{\bar{R}} (1 - H(s))^{k-1} ds}{\sum_{k=1}^I q_k k (1 - H(r))^{k-1}}. \quad [2.1]$$

It is easy to verify that $\beta(\cdot; \mathbf{q})$ is non-decreasing. In particular $\beta(\cdot; \mathbf{q})$ is continuously differentiable with the following derivative

$$\beta'(r; \mathbf{q}) = \frac{h(r) \left(\sum_{k=2}^I q_k k (k-1) (1 - H(r))^{k-2} \right) \left(\sum_{k=1}^I q_k k \int_{s=r}^{\bar{R}} (1 - H(s))^{k-1} ds \right)}{\left(\sum_{k=1}^I q_k k (1 - H(r))^{k-1} \right)^2}, \quad [2.2]$$

where $h(\cdot)$ denotes the probability density function (PDF) of R_i . The form of the derivative suggests that: if $q_1 = 1$ then $\beta'(r; \mathbf{q}) = 0$ for all r ; otherwise $\beta(\cdot; \mathbf{q})$ will be strictly increasing almost everywhere. We shall focus on the latter case as $\beta(R_i; \mathbf{q})$ has a continuous distribution.

2.2.2 Demand Side

All consumers have the same valuation of the object but differ in search cost. Each draws a search cost c from a continuous distribution with CDF $G(\cdot)$.⁸ She decides how many firms to visit before conducting the search. Then a consumer with search cost c faces the following decision problem:

$$\min_{k \geq 1} c(k-1) + \mathbb{E}_F [P_{(1:k)}].$$

We use $\mathbb{E}_F[\cdot]$ to denote an expectation where the random prices have distribution described by the CDF $F(\cdot)$. As standard, we assume there is no cost for the first search. The valuation of the object is set to \bar{R} , so that purchase is always made.

⁸To ease notation we suppress any consumer-specific subscripts from the exposition of the model.

Consumer's Best Response

It is easy to verify that $\mathbb{E}_F [P_{(1:k)}]$ is non-increasing in k , and we have strict monotonicity when price has a non-degenerate distribution. The marginal saving from searching one more store after having searched k stores is:

$$\Delta_k (F) \equiv E_F [P_{(1:k)}] - E_F [P_{(1:k+1)}].$$

$\Delta_k (F)$ is also non-increasing in k . When price has a continuous distribution it can be shown that

$$\Delta_k (F) = \int F (p) (1 - F (p))^k dp. \quad [2.3]$$

It then follows that the proportions of consumers searching optimally will satisfy this rule:

$$q_k (F) = \begin{cases} 1 - G (\Delta_1 (F)) & \text{for } k = 1 \\ G (\Delta_{k-1} (F)) - G (\Delta_k (F)) & \text{for } k > 1 \end{cases}. \quad [2.4]$$

Since the search behaviour in our model is standard, the same expression is to be found in Hong and Shum (2006) and Moraga-González and Wildenbeest (2007).

2.2.3 Equilibrium

For any $\mathbf{q} \in \mathbb{S}^{I-1}$, $\beta (\cdot; \mathbf{q})$ in (2.1) gives an expression for the firm's best response that induces a price distribution. Conversely, given any price CDF, $F (\cdot)$, (2.4) gives the consumer's best response $\mathbf{q} (F) = (q_k (F))_{k=1}^I$. Therefore we can define a symmetric equilibrium for our game as follows.

DEFINITION (Symmetric Bayesian Nash equilibrium). *The pair $(\mathbf{q}, \beta (\cdot; \mathbf{q}))$ is a symmetric equilibrium if:*

(i) *for every \mathbf{q} when all firms apart from i use pricing strategy $\beta (\cdot; \mathbf{q})$, $\beta (\cdot; \mathbf{q})$ is a best response for firm i ;*

(ii) *given the price distribution induced by $\beta (\cdot; \mathbf{q})$, \mathbf{q} is a vector of proportions of consumers' optimal search.*

For example the monopoly pricing strategy when all consumers search just once constitutes an equilibrium with: $\beta^M (r; \mathbf{q}^M) = \bar{R}$ for all r , and \mathbf{q}^M such that $q_1 = 1$ and $q_k = 0$ for $k \neq 1$. However, $(\mathbf{q}^M, \beta^M (\cdot; \mathbf{q}^M))$ does not generate any price dispersion and can thus be easily refuted by the data. We will focus on an equilibrium where consumers search more than once with a positive measure. Since $\beta (\cdot; \mathbf{q})$ gives firms' best response given any search behaviour, the

equilibrium can be characterised by \mathbf{q} that satisfies (2.1) and (2.4) simultaneously. We state this as a proposition.

PROPOSITION 1. *In an equilibrium with strictly increasing pricing strategy with an inverse function $\xi(\cdot; \mathbf{q})$, \mathbf{q} satisfies the following system of equations:*

$$q_k = \begin{cases} 1 - G\left(\int H(\xi(p, \mathbf{q})) (1 - H(\xi(p, \mathbf{q}))) dp\right) & \text{for } k = 1 \\ G\left(\int H(\xi(p, \mathbf{q})) (1 - H(\xi(p, \mathbf{q})))^k dp\right) - G\left(\int H(\xi(p, \mathbf{q})) (1 - H(\xi(p, \mathbf{q})))^{k+1} dp\right) & \text{otherwise} \end{cases} \quad [2.5]$$

The characterisation above states that an equilibrium can be summarised by a fixed-point of some map, say \mathcal{T} . It can be shown using the implicit function theorem that \mathcal{T} is a continuous map under some regularity conditions.⁹ It is clear that \mathcal{T} maps \mathbb{S}^{I-1} to some subset of \mathbb{S}^{I-1} . Therefore a general proof for an existence of an equilibrium with a price dispersion may be shown by using a fixed-point theorem, such as Brouwer's, by showing that \mathcal{T} maps a certain subset of \mathbb{S}^{I-1} onto itself.

In subsequent sections we shall assume an existence of an equilibrium characterised by Proposition 1. We henceforth drop the indexing arguments of equilibrium objects that are made explicit in this Section for the purpose of discussions on best response; e.g. $\beta(\cdot; \mathbf{q})$ becomes $\beta(\cdot)$, $\mathbb{E}_F[\cdot]$ becomes $\mathbb{E}[\cdot]$ etc.

2.3 Identification

We identify the demand side first then proceed to the supply side. Our identification of the demand side focuses on \mathbf{q} . We assume another variable related to price is available. Once we can identify \mathbf{q} , identification of the firm's cost distribution follows analogously to GPV.

2.3.1 Demand Side

Suppose we know the equilibrium price distribution of a search model. This is expected if we have a random sample $\{P_{im}\}_{i=1, m=1}^{I, M}$ of prices for I firms from M markets, and we let $M \rightarrow \infty$. By assumption $P_{im} = \beta(R_{im})$. Here Y_{im} denotes an observable variable that is assumed to satisfy

⁹Continuity makes the application of Brouwer's theorem straightforward, as opposed to the existence proof in Moraga-González, Sándor, and Wildenbeest (2017) where the mapping needed to be transformed to address the discontinuity around 0.

Assumption I below. The main identifying assumption we introduce in our work links Y_{im} to the expected probability firm i winning the sale of the object conditioning on setting price to be P_{im} .

ASSUMPTION I. *There exists a finite and positive λ such that*

$$\mathbb{E}[Y_{im}|P_{im}] = \lambda \sum_{k=1}^I q_k \frac{k}{I} (1 - F(P_{im}))^{k-1}. \quad [2.6]$$

The expression above says: Y_{im} is proportional to the probability firm i wins with price P_{im} . Assumption I is analogous to the well-known assumption in the demand estimation literature in IO that equates the observed market share with the choice probabilities; e.g. as used in Berry, Levinsohn, and Pakes (1995). In our case, depending on the context, candidates for Y_{im} could be market share or sales volume. The unknown λ does not prevent identification since we have the restriction that $\sum_{k=1}^I q_k$ must be 1. It is important to note that unlike in a discrete choice model, where the choice probabilities sums to 1, the ex-post probability $\sum_{k=1}^I q_k \frac{k}{I} (1 - F(P_{im}))^{k-1}$ will almost surely not sum to one across i . The role of λ in equation (2.6) ensures \mathbf{q} can be interpreted independently from this scale. For simplicity we assume λ to be the same for all m but this is not necessary.

Let $Y_m = (Y_{1m}, \dots, Y_{Im})^\top$ and \mathbf{X}_m be an $I \times I$ matrix such that $(\mathbf{X}_m)_{ik} = \frac{k}{I} (1 - F(P_{im}))^{k-1}$. We vectorize Y_m and \mathbf{X}_m across m to form: $\mathbf{Y} = [\mathbf{Y}_1^\top : \dots : \mathbf{Y}_M^\top]^\top$ and $\mathbf{X} = [\mathbf{X}_1^\top : \dots : \mathbf{X}_M^\top]^\top$. Then under Assumption I, we have

$$\mathbf{q} = \frac{\mathbb{E}[\mathbf{X}^\top \mathbf{X}]^{-1} \mathbb{E}[\mathbf{X}^\top \mathbf{Y}]}{\iota^\top \mathbb{E}[\mathbf{X}^\top \mathbf{X}]^{-1} \mathbb{E}[\mathbf{X}^\top \mathbf{Y}]}, \quad [2.7]$$

where ι denotes a $IM \times 1$ vector of ones. Note that \mathbf{X} has full rank almost surely when P_{im} has a continuous distribution as columns in \mathbf{X}_m form a polynomial basis of $\left\{ (1 - F(P_{im}))^{l-1} \right\}_{l=1}^I$. Generally \mathbf{q} is overidentified in the sense that it can be identified using $(\mathbf{Y}_m, \mathbf{X}_m)$ for any m when $F(\cdot)$ is known.

2.3.2 Supply Side

The optimal strategy derived in (2.1) expresses the optimal price in terms of the latent marginal cost. Although such expression is intuitive and natural from the theoretical analysis, it is not immediately useful for empirical purposes.¹⁰ We instead consider defining $\beta(\cdot)$ as a maximiser of

¹⁰It is, however, useful for generating data in simulation studies.

the following function:

$$\Lambda(p, r) = (p - r) \sum_{k=1}^I q_k \frac{k}{I} (1 - H(\xi(p)))^{k-1}.$$

Taking a (partial) derivative of the above with respect to p gives,

$$\begin{aligned} \frac{\partial}{\partial p} \Lambda(p, r) &= \sum_{k=1}^I q_k \frac{k}{I} (1 - H(\xi(p)))^{k-1} \\ &\quad + (p - r) \xi'(p) h(\xi(p)) \sum_{k=1}^I q_k \frac{k(k-1)}{I} (1 - H(\xi(p)))^{k-2}. \end{aligned}$$

We next use the insight from GPV by relating the distributions between the observed and unobserved variables. Particularly:

$$F(p) = H(\xi(p)) \quad \text{and} \quad f(p) = \xi'(p) h(\xi(p)),$$

so that the first order condition implies

$$\sum_{k=1}^I q_k k (1 - F(p))^{k-1} = (p - \xi(p)) f(p) \sum_{k=2}^I q_k k (k-1) (1 - F(p))^{k-2}.$$

We then obtain the explicit form for $\beta^{-1}(\cdot)$ as,

$$\xi(p) = p - \frac{\sum_{k=1}^I q_k k (1 - F(p))^{k-1}}{f(p) \sum_{k=2}^I q_k k (k-1) (1 - F(p))^{k-2}}. \quad [2.8]$$

We can identify R_i from P_i , $f(\cdot)$, $F(\cdot)$ and $\{q_k\}_{k=1}^I$. Thus we can identify $\{R_i\}_{i=1}^I$ through $\{\xi(P_i)\}_{i=1}^I$, and subsequently identify $h(\cdot)$ with data from multiple markets.

2.3.3 Constructive Identification

Suppose we have a random sample for firms from multiple markets $\{(P_{im}, Y_{im})\}_{i=1, m=1}^{I, M}$. There is a natural corresponding estimation strategy by replacing unknown population quantities by sample analogs.

Estimation of \mathbf{q}

We first construct an estimator for $F(\cdot)$, such as the empirical CDF. We can estimate \mathbf{q} using the sample counterpart of (2.7); by removing the expectation operators and replace \mathbf{X} by its estimate $\widehat{\mathbf{X}}$ that replaces the unknown $F(\cdot)$ by some estimator $\widehat{F}(\cdot)$. Then

$$\widehat{\mathbf{q}} = \frac{\left(\widehat{\mathbf{X}}^\top \widehat{\mathbf{X}}\right)^{-1} \widehat{\mathbf{X}}^\top \mathbf{Y}}{\iota^\top \left(\widehat{\mathbf{X}}^\top \widehat{\mathbf{X}}\right)^{-1} \widehat{\mathbf{X}}^\top \mathbf{Y}}.$$

Our estimator of \mathbf{q} is a smooth functional of an estimator of $F(\cdot)$. Therefore $\widehat{\mathbf{q}}$ is expected to converge at the parametric rate of \sqrt{M} .

Estimation of $h(\cdot)$

We start by obtaining an estimate for R_{im} using:

$$\widehat{R}_{im} = P_{im} - \frac{\sum_{k=1}^I \widehat{q}_k k \left(1 - \widehat{F}(P_{im})\right)^{k-1}}{\widehat{f}(P_{im}) \sum_{k=1}^I \widehat{q}_k k (k-1) \left(1 - \widehat{F}(P_{im})\right)^{k-2}}, \quad [2.9]$$

here $\widehat{f}(\cdot)$ and $\widehat{F}(\cdot)$ are some estimators for $f(\cdot)$ and $F(\cdot)$ respectively. We can then perform nonparametric density estimation for $h(\cdot)$ with $\left\{\widehat{R}_{im}\right\}_{i=1, m=1}^{I, M}$. When we estimate $f(\cdot)$ and $F(\cdot)$ nonparametrically, it is expected that the rate of convergence of \widehat{R}_{im} (and subsequently the estimator of $h(\cdot)$) will be determined by $\widehat{f}(\cdot)$; both $\widehat{\mathbf{q}}$ and $\widehat{F}(\cdot)$ converge at a faster rate.

2.4 Main Results

In the following section we present two theorems. Theorem 1 shows that the theoretical search model imposes testable restrictions on the distribution of the observed prices. Theorem 2 gives the convergence rate for $\widehat{h}(\cdot)$.

2.4.1 Nonparametric Restrictions on the Data

Let \mathcal{P} denote the set of strictly increasing CDFs with support in \mathbb{R} . Let $\mathbf{F}(\cdot)$ denote the joint CDF of equilibrium prices.

THEOREM 1. *Let $I \geq 2$. Let $\mathbf{F}(\cdot) \in \mathcal{P}^I$ with support $[\underline{P}, \overline{P}]^I$. There exists a distribution of marginal cost with CDF $H(\cdot)$, with an increasing CDF $H(\cdot) \in \mathcal{P}$ such that $\mathbf{F}(\cdot)$ is the joint CDF of the equilibrium prices in the search model if and only if:*

$$C1. \mathbf{F}(p_1, \dots, p_K) = \prod_{i=1}^I F(p_i);$$

C2. $\xi(\cdot)$ defined in (2.8) is strictly increasing on $[\underline{P}, \overline{P}]$, and its inverse is differentiable on $[\underline{R}, \overline{R}] = [\xi(\underline{P}), \xi(\overline{P})]$.

Moreover, when $H(\cdot)$ exists, it is unique with support $[\underline{R}, \overline{R}]$ and satisfies $H(r) = F(\xi^{-1}(r))$ for all $r \in [\underline{R}, \overline{R}]$. In addition, $\xi(\cdot)$ is the quasi-inverse of the equilibrium strategy in the sense that $\xi(p) = \beta^{-1}(p)$ for all $r \in [\underline{P}, \overline{P}]$.

Our Theorem 1 is analogous to Theorem 1 in GPV and the proof is presented in appendix B.1.

2.4.2 Large Sample Properties

In order to study the rate of convergence of our estimators we need to know some regularity properties of the objects to be estimated. We begin with some regularity assumptions on the distribution of the underlying cost.

ASSUMPTION A.

(i) For any observed price P : there exists R such that

$$P = R + \frac{\sum_{k=1}^I q_k k \int_{s=R}^{\overline{R}} (1 - H(s))^{k-1} ds}{\sum_{k=1}^I q_k k (1 - H(R))^{k-1}},$$

for \mathbf{q} that satisfies Proposition 1, and there is an observable Y that satisfies Assumption I;

(ii) $H(\cdot)$ admits upto $\tau + 1$ continuous derivatives on $[\underline{R}, \overline{R}]$.

The equilibrium restrictions imply the following properties for the observed price distribution:

PROPOSITION 2. Under Assumption A:

$$(i) f(p) = \frac{1}{p - \xi(p)} \left(\frac{\sum_{k=1}^I q_k k (1 - F(p))^{k-1}}{\sum_{k=1}^I q_k k (k-1) (1 - F(p))^{k-2}} \right);$$

(ii) $\inf_{p \in [\underline{P}, \overline{P}]} f(p) > 0$;

(iii) $\lim_{p \rightarrow \overline{P}} f(p) = \infty$, furthermore $0 < \lim_{p \rightarrow \overline{P}} \frac{f(p)}{(\overline{P} - p)^{-1}} < \infty$;

(iv) $F(\cdot)$ admits upto $\tau + 1$ continuous derivatives on $[\underline{P}, \overline{P}]$;

(v) $f(\cdot)$ admits upto $\tau + 1$ continuous derivatives on $(\underline{P}, \overline{P})$.

The findings we want to highlight here are (iii) and (v). The former reveals that $f(\cdot)$ has a pole at the upper boundary. Kernel density estimation in a neighbourhood of a pole has to be treated with care (e.g. see Section 5 in Marron and Ruppert (1994)). We suggest a transformation to deal with this issue below.¹¹ The latter suggests that the implied observed PDF is smoother than the latent PDF; similar findings are also found in GPV based on the same rationale by an inspection of (i).

Suppose we have data $\{(P_{im}, Y_{im})\}_{i=1, m=1}^{I, M}$. We assume to have some preliminary estimators for \mathbf{q} , $F(\cdot)$, and $f(\cdot)$ that converge to zero at some rates as $M \rightarrow \infty$. Let $\eta_{0, M} = \left(\frac{\log M}{M}\right)^{\frac{\tau+1}{2\tau+3}}$, so that $\eta_{0, M}$ is the optimal rate of convergence for density estimation with $\tau + 1$ continuous derivatives (see Stone (1982)).

ASSUMPTION B. *Suppose $\{(P_{im}, Y_{im})\}_{i=1, m=1}^{I, M}$ satisfies Assumption A. There exists estimators: $\hat{\mathbf{q}}$, $\hat{F}(\cdot)$, and $\hat{f}(\cdot)$ such that:*

- (i) $\|\hat{\mathbf{q}} - \mathbf{q}\| = O\left(1/\sqrt{M}\right)$ a.s.;
- (ii) $\sup_{p \in [\underline{P}, \bar{P}]} \left| \hat{F}(p) - F(p) \right| = O\left(1/\sqrt{M}\right)$ a.s.;
- (iii) *For any positive sequence ε'_M that decreases to 0 there exists some positive sequence δ'_M that decreases to zero such that $\sup_{p \in [\underline{P} + \delta'_M, \bar{P} - \delta'_M]} \left| \hat{f}(p) - f(p) \right| = o\left(\frac{\eta_{0, M}}{\varepsilon'_M}\right)$ a.s.;*
- (iv) *There exist some positive sequences $\{\delta_M\}$ and $\{\eta_M\}$ that decrease to zero such that $\eta_{0, M} = o(\eta_M)$, $\sup_{p \in [\underline{P} + \delta_M, \bar{P} - \delta_M]} \left| \hat{f}(p) - f(p) \right| = O(\eta_M)$ a.s.*

Estimators for \mathbf{q} and $F(\cdot)$ that converge at a parametric rate are going to be available under weak conditions. We will focus on the uniform convergence properties of a kernel estimator for $\hat{f}(\cdot)$. Studying uniformity over the entire support of P_{im} is difficult as the support is compact. It is well-known that kernel estimators have problems at (and near) the boundaries; e.g. see Chapter 2.11 in Wand and Jones (1999). On the other hand, if we consider any fixed inner subset of $[\underline{P}, \bar{P}]$, then a kernel density estimator can achieve the convergence rate $\eta_{0, M}$ under standard constructions, for example by using a $\tau + 1$ order kernel and setting the bandwidth to be proportional to $b_{0, M} \equiv \left(\frac{\log M}{M}\right)^{\frac{1}{2\tau+3}}$; see Härdle (1991). But these rates cannot be maintained when we allow the support to expand to $[\underline{P}, \bar{P}]$ as sample size grows. Existing results on the uniform convergence rates for kernel estimators over expanding supports assume densities are bounded (e.g. see Masry

¹¹There are also other auction models that have unbounded densities. E.g. in a first price auction with a reserve price (see GPV) and in models with selective entry (see Gentry, Li, and Lu (2015)).

(1996) and Hansen (2008)). They are therefore not immediately applicable to us due to the pole at \bar{P} .

Assumption B.(iii) says that any decreasing function of M converging to zero slower than $\eta_{0,M}$ can serve as an *upper bound* for $\sup_{p \in [\underline{P} + \delta'_M, \bar{P} - \delta'_M]} \left| \hat{f}(p) - f(p) \right|$ for some $\delta'_M = o(1)$. This is possible, for instance, with a kernel estimator using a transformation method. From Proposition 2.(iii) we know $f(p)$ behaves similarly to $(\bar{P} - p)^{-1}$ for p close to \bar{P} . Then let us consider $P_{im}^\dagger \equiv -\ln(\bar{P} - P_{im})$. The support of P_{im}^\dagger is $[-\ln(\bar{P} - \underline{P}), \infty)$. Denote the PDF of P_{im}^\dagger by $f^\dagger(\cdot)$. By a change of variable, we have:

$$f(p) = \frac{f^\dagger(-\ln(\bar{P} - p))}{\bar{P} - p}.$$

It then follows that $f^\dagger(\cdot)$ is bounded and, in particular, $f^\dagger(-\ln(\bar{P} - p))$ is flat as $p \rightarrow \bar{P}$.

Furthermore it has the same smoothness as $f(\cdot)$,¹² Consider the following estimators:

$$\begin{aligned} \hat{f}(p) &= \frac{\hat{f}^\dagger(-\ln(\bar{P} - p))}{\bar{P} - p}, \text{ where} \\ \hat{f}^\dagger(p^\dagger) &= \frac{1}{MIb_M^\dagger} \sum_{m=1}^M \sum_{i=1}^I K\left(\frac{P_{im}^\dagger - p^\dagger}{b_M^\dagger}\right) \text{ for any } p^\dagger, \end{aligned}$$

and $K(\cdot)$ is a kernel function with a bandwidth b_M^\dagger . Thus it can be shown that $\hat{f}^\dagger(\cdot)$ converges uniformly at rate $\eta_{0,M}$ over some expanding support when we use a $\tau + 1$ higher order kernel coupled with bandwidth $b_{0,M}$. The division by $\bar{P} - p$ slows down the rate of convergence for $\hat{f}(\cdot)$ at the upper boundary. This can be controlled to be as slow as we like by letting δ'_M go to zero slowly. There is also a bias issue at the lower boundary. This can be avoided by setting $b_M^\dagger = o(\delta'_M)$.

Assumption B.(iv) then assumes an existence of an estimator for $f(\cdot)$ that converges uniformly over $[\underline{P} + \delta_M, \bar{P} - \delta_M]$ at an *achievable rate* η_M . We can extend the argument given for B.(iii) and make η_M arbitrarily close to $\eta_{0,M}$. More specifically, we can set $\delta_M = \bar{P} - \varepsilon_M$ for some decreasing positive sequence $\{\varepsilon_M\}$ such that $b_{0,M} = o(\delta_M)$. Then B.(iv) holds with $\eta_M = \frac{\eta_{0,M}}{\varepsilon_M}$.

Now that we have some estimators that satisfy Assumption B, we turn to \hat{R}_{im} as defined in equation (2.9). We shall use a modified version of \hat{R}_{im} for the second stage estimation since we only have the desired uniform convergence rate for $\hat{f}(\cdot)$ over an expanding support. For some

¹²For any $p^\dagger \in [-\ln(\bar{P} - \underline{P}), \infty)$, $f^\dagger(p^\dagger) = \exp(-p^\dagger) f(\bar{P} - \exp(-p^\dagger))$.

positive sequence $\{\delta_M\}$ that decrease to zero, let

$$\tilde{R}_{im} = \begin{cases} \hat{R}_{im} & \text{for } P_{im} \in [\underline{P} + \delta_M, \bar{P} - \delta_M] \\ +\infty & \text{otherwise} \end{cases}. \quad [2.10]$$

When $\tilde{R}_{im} < \infty$, \tilde{R}_{im} is a smooth function of $\hat{\mathbf{q}}$, $\hat{F}(\cdot)$ and $\hat{f}(\cdot)$. Therefore we can obtain its convergence rate that is determined by $\sup_{p \in [\underline{P} + \delta_M, \bar{P} - \delta_M]} |\hat{f}(p) - f(p)|$.

LEMMA 1. *Under Assumptions A and B, for the same $\{\delta_M\}$ and $\{\eta_M\}$ in B(d),*

$$\sup_{i,m \text{ s.t. } \tilde{R}_{im} < \infty} |\tilde{R}_{im} - R_{im}| = O(\eta_M) \text{ a.s.}$$

We explicitly define a kernel estimator for $h(\cdot)$ here as:

$$\hat{h}(r) = \frac{1}{MIb_M} \sum_{m=1}^M \sum_{i=1}^I K\left(\frac{\tilde{R}_{im} - r}{b_M}\right) \text{ for any } r.$$

As before, $K(\cdot)$ is a kernel function with a bandwidth b_M . We can use Lemma 1 to quantify the estimation error that arises from using \tilde{R}_{im} instead of R_{im} , and obtain the convergence rate for $\hat{h}(\cdot)$.

THEOREM 2. *Under Assumptions A and B, and for the same $\{\delta_M\}$ and $\{\eta_M\}$ as in B.(iv), let: (i) $K(\cdot)$ be a symmetric $(\tau + 1)$ -order kernel with support $[-1, 1]$; (ii) $K(\cdot)$ is twice continuously differentiable on $[-1, 1]$; (iii) $\{b_M\}$ for some positive real numbers decreasing to zero such that $\delta_M = O(b_M)$. Then for any sequence $\{\varsigma_M\}$ of positive real numbers decreasing to zero such that $b_M = o(\varsigma_M)$,*

$$\sup_{r \in [\underline{R} + \varsigma_M, \bar{R} - \varsigma_M]} |\hat{h}(r) - h(r)| = O\left(\frac{\eta_M}{b_M}\right) \text{ a.s.}$$

Theorem 2 shows that $\hat{h}(\cdot)$ converges at a slower rate than $\hat{f}(\cdot)$ by a factor of b_M^{-1} . We have argued that the convergence rate for the latter can be made arbitrarily close to $\eta_{0,M}$. Therefore an appropriate choice of b_M will ensure $\hat{h}(\cdot)$ converges uniformly at a rate arbitrarily close to $\frac{\eta_{0,M}}{b_{0M}} = \left(\frac{\log M}{M}\right)^{\frac{\tau}{2\tau+3}}$, which is the optimal rate of convergence for a related density function derived in Theorem 3 of GPV.

2.5 Possible Extensions

We briefly discuss how to extend our model and methodology. First we generalise Assumption I by allowing for possibly nonparametric relation between Y_i , and the probability that firm i wins the sale with price P_i . Then we consider an asymmetric game where firms have different probabilities of being found.

2.5.1 Relaxing Assumption I

We anticipate that Assumption I will be the most convenient in applications. However, the mathematical structure of the search problem is conducive to a nonparametric generalisation. In what follows let \mathbf{x}_{im} be a $I \times 1$ vector such that $(\mathbf{x}_{im})_k = \frac{k}{I} (1 - F(P_{im}))^{k-1}$.¹³

ASSUMPTION I'. *There exists a function $\phi : R \rightarrow R$ such that*

$$\mathbb{E}[Y_{im}|P_{im}] = \phi(\mathbf{x}_{im}^\top \mathbf{q}). \quad [2.11]$$

Assumption I is a parametric special case of Assumption I' when $\phi(\cdot)$ is an identity function multiplied by an unknown scale. More generally Assumption I' only imposes that: Y_i is a (possibly unknown) function of the probability firm i wins with price P_{im} . When $\phi(\cdot)$ is parametrically specified, whether \mathbf{q} is identifiable depends on the parametric specification. A sufficient, but not necessary, condition for identification is strict monotonicity of $\phi(\cdot)$.¹⁴ When $\phi(\cdot)$ is unknown, (2.11) imposes a semiparametric index restriction. Ichimura (1993, Theorem 4.1) provides a set of conditions for identification of an index model like ours. Note that we cannot apply, at least without any modification, the average derivative argument of Powell, Stock, and Stoker (1989) to identify \mathbf{q} as our model does not satisfy their boundary conditions (see their Assumption 2). When \mathbf{q} is identified, regardless whether $\phi(\cdot)$ is known or not, we would expect the estimator for \mathbf{q} to converge sufficiently fast to not affect the convergence rate for $\hat{f}(\cdot)$, and subsequently $\hat{h}(\cdot)$ under general conditions.

¹³In principle we can also allow \mathbf{w}_{im} to be other known functions of $\{P_{im}\}_{i=1}^I$. But \mathbf{q} has a structural meaning so it is natural to use powers of the price hazard functions as in Assumption I.

¹⁴Let $\phi^{-1}(\cdot)$ denote the inverse of $\phi(\cdot)$. Given that $\mathbb{E}[\mathbf{x}_{im}\mathbf{x}_{im}^\top]$ has full rank a sufficient we can write $\mathbf{x}_{im}^\top \mathbf{q} = \phi^{-1}(\mathbb{E}[Y_{im}|P_{im}])$, so that

$$\mathbf{q} = \frac{\mathbb{E}[\mathbf{x}_{im}\mathbf{x}_{im}^\top]^{-1} \mathbb{E}[\mathbf{x}_{im}\phi^{-1}(\mathbb{E}[Y_{im}|P_{im}])]}{I^\top \mathbb{E}[\mathbf{x}_{im}\mathbf{x}_{im}^\top]^{-1} \mathbb{E}[\mathbf{x}_{im}\phi^{-1}(\mathbb{E}[Y_{im}|P_{im}])]}.$$

2.5.2 Asymmetric Search Probabilities

Consider a situation when firms have different probabilities of being searched. When a consumer sets out to visit k firms, for $\ell_i \in \{1, \dots, I\}$, we denote the probability that the set of firms $\{\ell_1, \dots, \ell_k\}$ get visited by $\omega_{\ell_1 \dots \ell_k}$. Since there is no need to keep track of different permutations of the same combination of firms, we only define $\omega_{\ell_1 \dots \ell_k}$ for $\ell_1 < \dots < \ell_k$. Let $\mathcal{I}_k \equiv \{\{\ell_1, \dots, \ell_k\} : \ell_j \in \mathcal{I} \text{ and } \ell_j < \ell_{j+1} \text{ for all } j\}$, and $\mathcal{I}_k^i \equiv \{\{\ell_1, \dots, \ell_k\} \in \mathcal{I}_k : \ell_j = i \text{ for some } j\}$. I.e. \mathcal{I}_k is the set of indices for all combinations of k firms. \mathcal{I}_k^i is the set of indices for all combinations of k firms that always include firm i . Let $C_k^I \equiv \frac{I!}{(I-k)!k!}$ denote the combinatorial number from choosing k objects from a set of I . Note that \mathcal{I}_k and \mathcal{I}_k^i have cardinality C_k^I and C_{k-1}^I respectively. Note that:

$$\omega_{\ell_1 \dots \ell_{i-1} \ell_{i+1} \dots \ell_k} = \sum_{\{\ell_1, \dots, \ell_k\} \in \mathcal{I}_k^i} \omega_{\ell_1 \dots \ell_k} \text{ for all } i, k.$$

Using a similar argument to the one applied previously, in equilibrium it can be shown that the optimal pricing strategy for firm i , $\beta_i(\cdot)$ becomes:

$$\beta_i(r) = r + \frac{\sum_{k=1}^I q_k \sum_{\{\ell_1, \dots, \ell_k\} \in \mathcal{I}_k^i} \omega_{\ell_1 \dots \ell_k} \int_{s=r}^{\bar{R}} \prod_{j:1 \leq j \leq k, \ell_j \neq i} (1 - H(\xi_{\ell_j}(\beta_i(s)))) ds}{\sum_{k=1}^I q_k \sum_{\{\ell_1, \dots, \ell_k\} \in \mathcal{I}_k^i} \omega_{\ell_1 \dots \ell_k} \prod_{j:1 \leq j \leq k, \ell_j \neq i} (1 - H(\xi_{\ell_j}(\beta_i(r))))},$$

over the region of r where $\beta_i(\cdot)$ is strictly increasing.¹⁵ Here $\xi_j(\cdot)$ denotes the inverse of $\beta_j(\cdot)$. It is clear that we have asymmetric pricing functions that have been induced by differing probabilities of being searched.

We can also write down the inverse function for p defined over the region where $\beta_i(\cdot)$ is strictly increasing,

$$\xi_i(p) = p - \frac{\sum_{k=1}^I q_k \sum_{\{\ell_1, \dots, \ell_k\} \in \mathcal{I}_k^i} \omega_{\ell_1 \dots \ell_k} \prod_{j:1 \leq j \leq k, \ell_j \neq i} (1 - F_{\ell_j}(p))}{\sum_{k=2}^I q_k \sum_{\{\ell_1, \dots, \ell_k\} \in \mathcal{I}_k^i} \omega_{\ell_1 \dots \ell_k} \left(\sum_{j:1 \leq j \leq k, \ell_j \neq i} f_{\ell_j}(p) \prod_{j':1 \leq j' \leq k, \ell_{j'} \neq i, \ell_{j'} \neq j} (1 - F_{\ell_{j'}}(p)) \right)}.$$

We can extend Assumption I (and I') accordingly and replicate our earlier identification strategies. Particularly, we will need:

$$\mathbb{E}[Y_i | P_i] = \sum_{k=1}^I q_k \sum_{\{\ell_1, \dots, \ell_k\} \in \mathcal{I}_k^i} \omega_{\ell_1 \dots \ell_k} \int \prod_{j:1 \leq j \leq k, \ell_j \neq i} (1 - F_{\ell_j}(p)) dF_i(p).$$

¹⁵The support of optimal prices now differ between (some) firms.

2.6 Simulation

We consider a simple design for a game of search with $I = 3$. Consumers draw search costs from a distribution with CDF $G(c) = \sqrt{c}$ for $c \in [0, 1]$. Firms draw marginal costs from a uniform distribution on $[0, 1]$. We use the system of equations in (2.5) to iteratively solve for the equilibrium of the game.


We generate the data by drawing prices from (2.1) with $\mathbf{q} = (0.7852, 0.0455, 0.1693)$. We generate Y_i according to Assumption I with $\lambda = 1$. We generate the data for 333 markets, so $IM = 999$. We follow the estimation strategy described in Section 3.3. In particular, we use the empirical price CDF to estimate $\hat{F}(\cdot)$. We employ different bias correction and transformation techniques to estimate the densities. For the bias correction we use the procedure proposed in Karunamuni and Zhang (2008, henceforth KZ) that has recently been shown to be effective when applied to auction models (see Hickman and Hubbard (2015), and Li and Liu (2015)). We use the Epanechnikov kernel along with the forms of the plug-in bandwidths suggested in KZ. Since KZ's technique does not accommodate unbounded densities we also use the transformation we suggested in Section 4.2 to address the upper support. We combine it with the KZ's estimator to correct for the bias at the lower boundary of the support. We repeat the experiment 10000 times. The parametric estimators work extremely well and thus we only show graphs for the potentially more problematic density estimation.

Figure 2.1 shows the true price density, the mean, 2.5th and 97.5th percentiles (dotted lines) of the boundary corrected kernel estimator of KZ (in blue) and the kernel estimator that transforms the data to deal with the pole (in red). It is clear that standard boundary correction procedure will not be sufficient to deal with unbounded densities. On the other hand the transformation method seems to serve the purpose very well.

We next consider three similar plots of the density estimation of marginal cost PDF (using KZ). In Figures 2.2–2.4, we use the true R_{im} to estimate the density (in blue) as the benchmark. The other density estimators (in red) in other figures contain estimated components. Those in Figures 2.2 and 2.3 are also infeasible as they estimate R_{im} using the unknown $f(\cdot)$: the former only estimates \mathbf{q} and the latter in addition estimates $F(\cdot)$. The result for the feasible estimator using \tilde{R}_{im} as defined in (2.10) is in Figure 2.4. Again, we plot the mean and the percentiles using solid and dotted lines, respectively.

Note that the boundary correction method of KZ does not completely eliminate the bias at the boundary even for the estimator that uses R_{im} . This is expected. We can, in fact, observe some improvements since density estimation without any bias correction would, in this case, converge


Figure 2.1: Monte Carlo performance of kernel density estimators for the distribution of prices.



Note: Red line shows the estimated density of prices using the transformation method for the upper limit of the support. Blue line is the boundary-corrected kernel density estimator of Karunamuni and Zhang (2008). Dotted lines correspond to the 2.5th and 97.5th percentiles in the simulations.

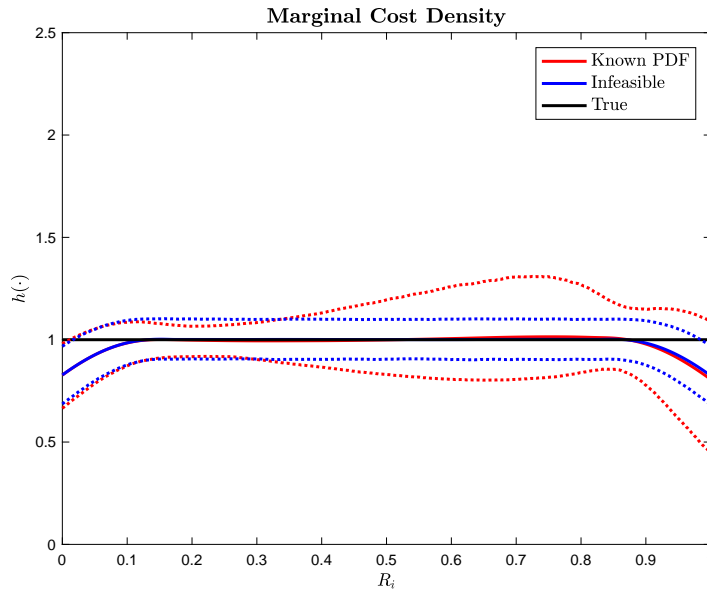
to 0.5 at both boundaries. The mean of the bandwidth use in these figures is around 0.17, and the estimator performs much better in the interior of the support away from the boundary by at least a bandwidth. Figures 2 - 4 also show that the main source of estimation error can be traced to the estimation of the price PDF. This is not unexpected given that the PDF is the most difficult object to estimate in the entire problem.

Figure 2.2: Monte Carlo comparison of two infeasible estimators.




Note: Red line shows the estimated density of marginal costs, assuming that the Blue line shows the performance of an infeasible estimator assuming that $\{F(\cdot), f(\cdot)\}$ are known instead of estimated. Blue line shows the performance of an infeasible estimator assuming that R_{im} are known. True density is $U[0, 1]$. Dotted lines correspond to the 2.5th and 97.5th percentiles in the simulations.

Figure 2.3: Monte Carlo comparison of two infeasible estimators.



Note: Red line shows the estimated density of marginal costs, assuming that the Blue line shows the performance of an infeasible estimator assuming that $\{f(\cdot)\}$ is known instead of estimated. Blue line shows the performance of an infeasible estimator assuming that R_{im} are known. True density is $U[0, 1]$. Dotted lines correspond to the 2.5th and 97.5th percentiles in the simulations.

Figure 2.4: Monte Carlo performance of the feasible density estimator.



Note: Red line shows the estimated density of marginal costs. Blue line shows the performance of an infeasible estimator assuming that R_{im} are known. True density is $U[0, 1]$. Dotted lines correspond to the 2.5th and 97.5th percentiles in the simulations.

2.7 Concluding Remarks

Hong and Shum (2006) and a series of papers by Moraga-González et al. show that we can identify the demand side of the market using just observed prices alone. We show when other market data,

such as market shares, are available we can allow firms to be heterogenous and identify the supply side as well.

We characterise the equilibrium in a search game with heterogenous consumers and firms that supports price dispersion. We provide conditions to identify the model and propose a way to estimate the model primitives. We show that the density of the unobserved marginal cost can be estimated to converge at an arbitrary close to, but not achieving, the optimal rate derived in related auction models (such as Guerre, Perrigne, and Vuong (2000)). The reason can be traced to the fact that the density of the equilibrium price has a pole at the upper support.

Chapter 3

Value of Information and the Impact of Mortgage Intermediaries on Lender Competition and Households' Financial Positions

3.1 Introduction

A large body of recent literature considers the role search frictions play in the determination of borrower choices in mortgage markets (Allen, Clark, and Houde (2013, 2017), Woodward and Hall (2012), Agarwal, Grigsby, Hortaçsu, Matvos, Seru, and Yao (2017), Alexandrov and Koulayev (2018), Guiso, Pozzi, Tsoy, Gambacorta, and Mistrulli (2017), Deltas and Li (2018)). Since a mortgage contract is the most complicated financial investment that most households undergo in their lifetime, the consequences of inadequate decisions can be serious. Searching for a mortgage involves not only comparing rates on different product types across banks, but also incurring the opportunity cost of filling out applications, being interviewed by lender representatives, and the risk that getting rejected might affect one's future credit score. The difficulty of the choice situation is further exacerbated by the pricing practices of lenders who use two-part tariffs and offer large menus of very similar products. For example, Kashyap and Rostom (2018) report that the median UK household taking out a mortgage faces a choice of 100 different options. No wonder, then, that many households choose to outsource this decision to mortgage intermediaries, or brokers, in the hopes of finding the best deal.

Differently from the US and Canada, the UK market is characterised by a very high percentage

of mortgages originated through brokers. The Intermediary Mortgage Lenders Association reports that in the second quarter of 2015, 67% of borrowers used broker services, which corresponded to 71% of total value of all new mortgages in that period (IMLA, 2015). This figure roughly corresponds to what we observe in our administrative data for 2016 and 2017. Since brokers in this market appear to be important agents affecting both borrower choices and lenders' pricing decisions, our work addresses two questions related to their activity: 1) what is the value of information they currently provide to consumers; 2) how would the market outcomes change if every mortgage applicant could costlessly access broker-like advice.

To answer these questions we formulate a structural model of search, in which heterogeneous consumers decide whether to use a broker or contact different lenders directly. Our choice of a search framework is motivated by a significant proportion of unexplained price dispersion we document in the data. On the supply side, lenders, who possess private information about the marginal costs of providing the loan, choose the price assuming that the same loan can be sold directly or through an intermediary.¹ Brokers are then treated as a platform which enables the borrowers who decide to use them to find the cheapest product, thereby reducing the monopoly power lenders exercise over poorly informed consumers. Under this assumption, the model becomes an extension of a search framework proposed by MacMinn (1980). Namely, the pricing problem is equivalent to a first-price procurement auction with an unknown number of competitors. This allows us to leverage the literature on nonparametric estimation of auction models and the identification results presented in the second chapter of this thesis to recover the unobserved primitives of the model – the distributions of borrower search costs and lender heterogeneity.

We use an extensive dataset on over 1.3M mortgage contracts from 2016 and 2017 to estimate these distributions conditional on a large set of consumer demographics and loan characteristics. Our structural estimates suggest that search cost distributions substantially differ across demographic groups with the median cost of obtaining an additional quote ranging from 5 to 30% of the median monthly mortgage cost of around £300. On average, search is more costly for borrowers who live in rural areas and non-first time buyers. On the supply side, we find that despite high concentration, the market is relatively competitive with an average Lerner index of 11.64%. Lenders' margins exhibit dispersion across mortgage types with higher LTV and longer term loans being on average less profitable.

To provide the answer to our first question, the estimates are then used to simulate optimal prices and search behaviour in a new equilibrium where intermediation is not available. We calcu-

¹Frankel (1998) called this *price coherence*, a term that was recently popularized in the economic literature by Edelman and Wright (2015).

late the value of information as the difference in the expected consumer surplus between the baseline and the counterfactual scenarios. The average effect of brokers for the entire market is positive and corresponds to savings of £112.15 per month on a median-sized mortgage. This figure can be further decomposed into savings due to lower prices (change of 33.7%) and lower average search expenditure (16.33% lower than without brokers). However, not all mortgagors benefit equally from the current market structure. Value of information is much higher for younger, low-income, first time buyers and all borrowers who choose 2-year fixed rate products. Remarkably, borrowers choosing longer fixed-rate deals or shorter amortisation periods could be better off if brokers were not present in the market. This finding can be linked to the low estimated dispersion of marginal costs, which in turn implies that marginal benefits from acquiring additional information are too low to justify paying broker commissions.

The overall positive effect can be attributed to the externality brokers impose on the direct market (Salz, 2017). The existence of intermediaries reduces lenders' market power who are unable to price discriminate between informed and uninformed consumers. This explanation is reinforced by looking at the counterfactual distribution of price-cost margins – with no intermediation, the average Lerner index reaches almost 28% and margins on over a fourth of all mortgages exceed 40%.

In a second counterfactual we look at the effects of a hypothetical market centralisation, which can be achieved by establishing a market-wide platform where lenders post prices and borrowers are automatically matched with the best offer. In the context of the UK market, this simulates the effects of growing popularity of online brokerage platforms, which not only work as price comparison tools, but rather act as online brokers and assist consumers through the entire application process. Under the assumption that direct sales are no longer possible and the platform is free to use by borrowers, we find that prices on a centralised market would on average decrease by 6.41% and the platform would allow consumers to save £27.63 per month. Borrowers, however, are the only agents in the market who benefit from centralisation in our model, as lenders' margins drop by nearly 50%. Since instead of human knowledge, online platforms use machine learning algorithms to generate advice, to provide a total welfare effect one would also need to weigh the rather modest reduction in prices and search expenditure against the the sunk cost of physical brokers exiting the market.

On the whole, this chapter makes two main contributions: to the best of our knowledge, we are the first to provide an estimate for the value of information added by mortgage brokers. We also document directly how changes in market structure affect pricing and competition between

banks. Since we allow for rich patterns of observed and unobserved heterogeneity, we are able to show that even though the net effects of brokers' presence are positive, not every borrower is better off in a world where intermediation is possible. Secondly, the structural model presented here is novel and provides an attractive framework for studying welfare effects in industries with two-sided platforms and search frictions. Importantly, the estimators do not require any optimisation, structural features are identified in closed form, and the results are robust to distributional assumptions about search costs and firm heterogeneity.

■ **Related literature.** This work is most closely related to the growing body of empirical papers using structural models of consumer search to study mortgage markets in countries other than the UK. Among those studies, Allen et al. (2017) is probably the one methodologically closest to our study, since the authors consider a search and bargaining framework with bilateral heterogeneity. The main difference is that Allen et al. (2017) focus on the role of loyalty advantage and do not study the role of intermediation, excluding brokered loans from the estimation sample. The paper by Woodward and Hall (2012) is on the opposite side of the spectrum, as it focuses solely on brokered mortgages, showing that mortgagors in the US market would be better off by obtaining offers from multiple brokers. We abstract from search for brokers and assume that the intermediaries operate in a competitive sector and have no incentives to provide dishonest advice. Guiso et al. (2017) do consider distorted advice and use Italian data to study whether in-house bank advisers steer borrowers into taking up more risky and expensive adjustable rate mortgages more frequently than fixed rate mortgages. Whereas they do find welfare losses associated with suboptimal advice, they also discover that banning advice altogether would result in an average annual loss of €998. This number is lower than our estimates, but since the advice in our model is considered to be fully impartial, we consider our estimates to be informative of the upper bound on the change in consumer surplus. In another recent study, Agarwal et al. (2017) use data on actual search behaviour and rejected mortgage applications to document that, contrary to predictions stemming from standard search models, more search does not always result in lower prices. To explain this finding, the authors introduce screening and the probability of getting one's mortgage application rejected into a standard search model, finding that a standard framework is only able to recover true search cost scaled by the probability of approval. While our data do not inform us about rejected applications, we remain agnostic whether our search cost estimates also indirectly account for the probability of being rejected. Alexandrov and Koulayev (2018) investigate the interplay of search and preference for non-price characteristics (such as brand effects) to explain suboptimal shopping efforts in the US market. Finally, even though Deltas and Li (2018) do not

have a structural model in their paper, they present an empirical evidence on how search costs in the US mortgage market can be reduced by network externalities.

We also contribute to the literature studying the role of brokers in retail financial markets (Bergstresser et al. (2007), Inderst and Ottaviani (2012b), Egan et al. (2018) Egan (2018)) and the relatively small literature using IO search models to study welfare effects of intermediation in other markets. For example, Gavazza (2016) investigates the role of dealers in the secondary market for business aircraft and Salz (2017) looks at the role of brokers in contracting trade waste removal in New York City. In particular, the structural model here resembles Salz's framework where the same firms participate in both direct and brokered markets and cannot charge different prices. Our main finding corroborates Salz's conclusion that overall, intermediation reduces information frictions and can be seen as a positive externality reducing market power. However, we are also able to show that the effects can be negligible or even negative for certain types of consumers. Our identification strategy relies on a weaker set of assumption and hence differs from Salz's approach. We discuss the econometric differences in detail in section 3.4.1.

Finally, this research is tangentially related to two strands of theoretical literature: an array of papers studying the effects of middlemen (e.g. Rubinstein and Wolinsky (1987), Biglaiser (1993), Yavaş (1994), Spulber (1995), Hall and Rust (2003)), and an active literature on multi-sided platforms (e.g. Armstrong (2006), Rochet and Tirole (2006), Galeotti and Moraga-González (2009), Edelman and Wright (2015), de Cornière and Taylor (2017)). The way we treat brokers in the model is reminiscent of a platform with endogenous buyer entry.

The chapter is organized as follows: section 3.2 outlines main institutional features of the industry, describes the data and provides some reduced-form evidence on price dispersion and the impact of brokers on transaction prices. Section 3.3 introduces the theoretical model and in section 3.4 we discuss nonparametric identification of its primitives and outline the estimation method. Our main results are presented in sections 3.5 and 3.6. Section 3.7 concludes and provides directions for future research.

3.2 The UK mortgage market and data

This section provides an overview of the main institutional features of the industry, describes the data used in our analysis and uses reduced-form techniques to establish several empirical facts about, which then guide the assumptions in the structural model.

3.2.1 Institutional overview

The UK mortgage market is relatively concentrated. Whilst there are a lot of financial institutions issuing mortgages, the concentration ratio of the six biggest banks exceeds 70%. As in the US, mortgage terms in the UK typically amortise over 25 years, although longer durations are also common². However, unlike the US, the contracts can be seen as short-term, as refinancing is common after the expiry of the initial period. The four most common products are 2-, 3-, and 5-year fixed rate mortgages (FRM) and 2-year adjustable rate mortgage (ARM). For our analysis, we focus only on FRMs, which make up over 90% of all mortgage contracts. Upon the expiration of the initial period, to keep the interest rate fixed, mortgagors can enter a new contract with the same or a different lender. The most popular type of mortgage in the UK is the 2-year fixed rate (over 60% of all contracts).

For each type of product, banks post quoted rates that vary according to the contract period and the ratio of the loan size to the property value (loan-to-value ratio, or LTV). When selecting a mortgage, and conditional on the LTV and loan size, borrowers face a trade-off: they can either choose to pay an upfront fee to the lender in return for lower monthly payments, or pay no fee but make higher monthly payments. Lender fees are small, however, especially relative to the size of the mortgage. Median loan fees are £999 and 40% pay no fees at all. As documented by Kashyap and Rostom (2018) and Iscenko (2018), among others, lenders typically offer broad product portfolios with different combinations of fees and interest rates. In addition to that, they also offer loans with optional cashback (a one-off lump sum payment to new borrowers) or flexible repayment schemes (i.e. possibility of over- or underpayment) which are priced differently.

A striking feature of the UK market is that almost 70% of mortgages are accessed via brokers. This number is significantly bigger than the share of brokered mortgages in the US (Alexandrov and Koulayev (2018) report roughly 10%) or Canada (Allen et al. (2017) have 28% of brokered contracts in their data). The *Intermediary Mortgage Lenders Association* (IMLA) report a general upward trend in the fraction of borrowers who use intermediaries by noting that since the financial crisis, the value share of mortgages originated via brokers has increased from about 50% to 71% in the second quarter of 2015 (IMLA, 2015). Applying for a mortgage directly typically involves face-to-face interviews at local bank branches, filling out lengthy application forms and facing the risk of rejection which can eventually impact one's credit rating. Brokers' task is provide advice about the most suitable product and assist borrowers through the application process. The market for intermediation is competitive and geographically dispersed: as noted by IMLA (2015), "(...)

²The median first time buyer amortises over 30 years.

the UK mortgage broking business is dominated by small firms serving local client bases. According to data from the Financial Adviser Confidence Tracking Index in September 2015, 69% of broking firms employed only 1 or 2 mortgage advisers with another 20% employing 3 to 5.”.

While there is no regulation in place that obliges all brokers to search through all available mortgage products³, broker services offer affordability comparisons across banks that are different from those of lenders’ in-house advisers. Intermediaries are compensated in one of three ways: they can receive commissions directly from borrowers, procurement fees from lenders, or charge both parties.⁴ In our data, the median fee paid by borrowers to brokers is relatively small: a lump sum of £349, or an average of £10 per month over the duration of the initial period.

3.2.2 Data

We use administrative data from the Product Sales Database (PSD), a loan-level dataset containing information on all new mortgage originations in the UK. The data contain information taken at the time of a mortgage application, including mortgagor characteristics such as age and income; loan details such as the issuing bank, interest rate, and loan size; and property details such as the purchase price and location. Whilst the data begin in 2005, the data quality is patchy until 2008. Following the financial crisis, data collection and monitoring substantially improved, and in 2016, key relevant variables were added, such as indicators for brokered/direct and fees. For this reason, our main analysis begins in 2016 and ends in the last quarter of 2017.

Our final sample includes over 1.3M contracts originated in the period of interest. Since we do not have data on any observable characteristics of brokers and do not know which of them search the entire market, we focus only on the big six lenders, assuming that their products are available to every potential borrower in the country. Another reason to focus on the big players only is because in our model we abstract from lenders’ budget constraints and capital requirements which are much more important for small lenders (Benetton, 2018). We also exclude adjustable rate mortgages which are a small fraction of the UK market and any loans with non-standard FRM period or LTV over 95%. Further details on the sample construction and summary statistics can be found in appendix C.1.1.

³The intermediaries that do that are known as *whole-of-market* brokers.

⁴Woodward and Hall (2012) argue that in the US brokers are indifferent between their source of compensation. A different strand of (mostly theoretical) literature studies how different compensation schemes can alter brokers’ incentives (see e.g. Inderst and Ottaviani (2012a)). Our study abstracts from this issue by assuming that any payments from lenders to brokers constitute a part of lenders’ costs which are eventually passed onto borrowers in the form of higher prices.

3.2.3 Mortgage cost

Since the total mortgage price faced by borrowers typically consists of two components (interest rate and upfront fees), in this sections we define a cost metric which allows us to compare different loans along one, unified dimension. Constructing a scalar measure of cost will turn out to be vital for the structural model, since all estimates will be denominated relative to it. To construct the measure, we use the fees, initial rate and initial fixed-rate deal period length to compute the monthly cost for the borrower. The monthly economic (“sunk”) cost is the interest component of the monthly payment plus any upfront fees added onto the loan by the lender:

$$p = iL + \frac{Fee}{N}, \quad [3.1]$$

where N is the initial period of the mortgage contract (2, 3, or 5 years), L is the size of the loan, and i is the fixed interest rate. Since in the structural model we take the loan size as given, to adequately compare costs of mortgages with different initial loan amounts, we normalise the monthly cost of the loan to correspond to a median loan value in the sample, £150,000. Our approach is similar to Allen et al. (2017) who normalise their price variable to correspond to the monthly payment on a \$100,000 loan.

3.2.4 Reduced form findings

In this section we use the PSD data to provide descriptive evidence of several features of the UK mortgage market that will lend credence to our choice of the modelling framework. First, we document dispersion in the transacted prices. Second, we show that borrowers who used brokers on average have lower monthly costs, but the sign changes once we factor in broker fees. Finally, we show that observable borrower and product characteristics are poor predictors for the decision whether to delegate search to a broker, which suggests that an underlying unobservable (such as search cost) might be the main factor driving this decision. Overall, the type of evidence we present is akin to that in section 3 of Salz (2017), which justifies using similar modelling assumptions.

Price dispersion

First we examine mortgage price dispersion by choice of sales channel (direct or broker). We do this by looking at the level of unexplained variation after regressing mortgage prices on observed characteristics, separately for the two different channels. More specifically, we run the following

hedonic regression:

$$p_{ijt} = \mathbf{X}'_{ijt}\boldsymbol{\beta} + \psi_t + \xi_j + u_{ijt} \quad [3.2]$$

where p_{ijt} is the mortgage price for household i , from bank j , at time t . In this analysis we use two different definitions of price – interest rate (where we control for the level of upfront fees on the right-hand side) and the normalised economic cost defined in section 3.2.3. \mathbf{X}_{ijt} is a vector of household and loan characteristics, e.g. household income, LTV, and the mortgage term.⁵ ψ_t and ξ_j are time and bank fixed effects⁶. Table 3.1 reports $1 - R^2$ (unexplained variation) and the coefficient of variation, using two different measures of mortgage cost for our dependent variable: the interest rate (in basis points) and normalised monthly interest payments (in £) during the initial period of the loan.

Panel A in Table 3.1 reports results using our first dependent variable, whilst Panel B reports results for our second measure. On average, the level of unexplained variation is about 30% and is higher for the outcome variable which includes upfront fees, but this varies by broker usage. This proportion is quantitatively similar to the percent of unexplained variation in the Canadian data reported by Allen et al. (2017) who report $1 - R^2$ of 0.39.⁷ The table also compares the results with and without lender fixed effects. Adding fixed effects to the specification allows us to control for any persistent differences in price across banks, whilst leaving the variation within banks and any transitory differences across lenders unexplained. Our results suggest that controlling for fixed effects substantially reduces the proportion of residual variation in the direct segment, but has virtually no effect on the R^2 in the regression using broker data. This finding is consistent with the assumption in our model that brokers help borrowers find the most suitable product across lenders. Suppose there exists a lender that, on average, sets higher interest rates than its competitors. Adding a fixed effect for that bank will help explain an additional portion of unobserved variation in prices in the direct segment of the market, especially if consumers do not shop around. However, in the brokered segment, intermediaries compare such a product against competing offers, so it is unlikely that borrowers who use a broker would be advised to take up such a product.

⁵More specifically, we control for household income, house price, loan size, LTV (included as a set of dummy variables corresponding to LTV thresholds), first-time buyer (FTB) status, region, mortgage type, length, and other product characteristics, as well as their interactions and allow for potential nonlinearities.

⁶We also run the same specification without bank fixed effects.

⁷See also Allen et al. (2014) for a detailed study of price dispersion in the Canadian market.

Table 3.1: Price dispersion by sales channel.

Panel A: Interest rate				
	No FE		With FE	
	Direct	Broker	Direct	Broker
$1 - R^2$	0.316	0.181	0.232	0.172
Coefficient of variation	0.307	0.316	0.307	0.316

Panel B: Interest payments				
	No FE		With FE	
	Direct	Broker	Direct	Broker
$1 - R^2$	0.369	0.364	0.283	0.355
Coefficient of variation	0.294	0.305	0.294	0.305

Note: Table presents $1-R^2$ from the regression defined by 3.2, separately by direct and broker sales channels and for two different definitions of price. The second row in each panel is the coefficient of variation defined the ratio of the standard deviation to the mean.

Price benefits provided by brokers

To establish whether brokered mortgages on average are cheaper, we check whether households who used a broker received a better price than households who did not after we control for a flexible function of individual and product characteristics. Table 3.2 reports results from regressions where the dependent variable is either the quoted interest rate (in basis points, bps) on a mortgage product or the monthly cost of the mortgage in interest and fees (our preferred definition from section 3.2.3). In all cases, the coefficient is negative and significant suggesting that those who shopped with a broker received a cheaper product. However, the monetary savings appear to be modest and are about 7 bps when measured by the interest rate and about £5 per month when measured in terms of monthly payments.⁸

Table 3.2: Price benefits of using a broker.

Dependent variable:	(1) Interest	(2) Interest	(3) Interest	(4) Monthly Payment	(5) Monthly Payment
Used a broker	-7.761*** (0.0824)	-6.710*** (0.0822)	-7.428*** (0.0816)	-5.103*** (0.119)	-3.562*** (0.122)
Lender Fees	Linear	Linear	Non-linear	-	-
Controls	Yes	Yes	Yes	Yes	Yes
Regional FE	No	Yes	Yes	No	Yes
Time FE	Yes	Yes	Yes	Yes	Yes
Observations	1,309,067	1,309,067	1,309,067	1,309,067	1,309,067
R^2	0.768	0.772	0.778	0.627	0.632

Note: *** denotes significant at 1% level. Robust standard errors in parentheses. Interest is measured in basis points. Monthly interest is the component of the initial monthly payment that goes towards payment of the interest, including lender fees, and normalised by the size of the loan. Controls are income, house price, loan size, LTV, first time buyer and mortgage term. Time fixed effects are at the monthly level. Regional fixed effects are at the Government Office Region level and include a flag for an urban region. Non-linearities in lender fees are controlled for using a fifth-order spline.

⁸Since some of the brokers in our sample are compensated directly by borrowers while others only receive commissions from the lenders, we test whether different broker compensation schemes affect their incentives to provide unbiased advice. We therefore run the same regression for two subsamples of the data – one which only includes brokers who are only paid by the lenders and one which only includes those who are not receiving any commissions from the banks. The sign and the magnitude of the effect measured by the coefficient of interest do not change by much across the subsamples, suggesting that brokers on average offer cheaper loans, regardless of who they are paid by. The results are presented in appendix C.1.4.

Table 3.3: Impact of using a broker on price plus broker fees.

Dependent variable:	(1)	(2)	(3)	(4)
	Monthly payment + broker fee			
Used a broker	1.729*** (0.122)	2.985*** (0.126)	7.049*** (0.140)	11.40*** (0.152)
Controls	Yes	Yes	Yes	Yes
Regional FE	No	Yes	No	Yes
Time FE	Yes	Yes	Yes	Yes
Observations	1,309,067	1,309,067	792,023	792,023
R^2	0.607	0.610	0.626	0.632

Note: *** denotes significant at 1% level. Robust standard errors in parentheses. Monthly interest is the component of the initial monthly payment that goes towards payment of the interest, including lender and broker fees, and normalised by the size of the loan. Columns (1) and (2) show estimates obtained using the entire sample whereas columns (3) and (4) restricts the sample to only brokers who charge borrowers directly (i.e. the fees are non-zero).

The dependent variable used in the regressions presented in table 3.2 is constructed in a way to control for lender fees only and our definition of monthly cost does not include broker commissions. Once we add broker fees divided by the number of months in the deal period to reflect monthly cost, the sign of the coefficient switches to positive (see table 3.3).

The findings summarised in tables 3.2 and 3.3 are in line with the descriptive evidence Salz (2017) used to justify the assumption that buyers with higher search cost select themselves into the brokered market. Brokers seem to be offering lower prices on average, but once their commissions are factored in, the final cost turns out to slightly higher than the average in the direct market. This finding is crucial to justify that borrowers with higher search costs are more likely to use brokers. Without the sign reversal, standard models of search would have difficulties explaining why brokers are not used by everyone in the market. To provide some intuition, suppose that one always expects to pay less by going to the broker. Then borrowers with low search costs would have an incentive to pretend that their cost is high and use them as well.

Predicting broker use

The final fact we document in this section is that observable characteristics of the borrower do not really help predict who uses brokers. Table C.1.6 in Appendix C.1.3 reports results from a linear probability model where we regressed the brokered/direct indicator on a large set of observable personal characteristics (e.g. age and income), mortgage product characteristics (e.g. product type and mortgage term), and regional fixed effects. Because of potential reverse causality issues once we control for product characteristics, we do not attempt to extrapolate our interpretation of the effects beyond conditional correlations. Irrespective of the observable characteristics that we control for, the R^2 is never greater than 0.13, even if we allow for multiple interactions between

variables.⁹ This is consistent with our hypothesis that observables have little predictive power in understanding who uses brokers. This opens up scope for an unobserved component, such as search cost, to be a more important driving force behind borrowers' decisions.

3.3 Model

This section introduces a stylised model of mortgage pricing when consumers can search across different lenders directly or use a broker. As in Allen, Clark, and Houde (2017) we assume that there exists an initial period outside the model where the borrower chooses the property she wishes to purchase, associated loan size, and the main characteristics of the mortgage, e.g. duration and whether she needs a flexible repayment scheme. Therefore, the dimension of search we consider is one where the borrower can compare similar products across different banks. The assumptions on the intermediation technology closely follow the ones in Salz (2017). We treat brokers as non-strategic players, and assume that they act in the borrowers' best interest by choosing the best offer available in the market at a given time.¹⁰ This assumption allows to treat brokers similarly to a price comparison platform, or, in the parlance of Baye et al. (2006), an *information clearinghouse*.¹¹

The framework we present here extends MacMinn (1980) and my search model with bilateral heterogeneity from chapter 2 by adding another stage to the consumers' problem where they decide whether to use an intermediary or search. Consider an environment with a finite number of J lenders and a continuum of borrowers with unit demands. Borrowers, indexed by i , receive iid draws from a continuous search cost distribution $\kappa_i \sim \mathcal{G}(\cdot|\mathbf{x}^G)$. \mathbf{x}^G is a vector of observables which can shift the distribution of search cost. They can be thought of covariates defining consumer type and include characteristics of the individuals, such as age, income, or location, as well as some characteristics that also describe the product (e.g. whether the borrower is a first-time buyer). Lenders are heterogenous in their marginal cost of providing the loan, $c_{ij} \sim \mathcal{H}(\cdot|\mathbf{x}^H)$, which is their private information. \mathcal{H} is continuously distributed on a compact support $[\underline{c}; \bar{c}]$.¹²

⁹Even though the R^2 is not a perfect measure of predictive power, we also looked at distributions of predicted probabilities (*propensity scores*) of using a broker for borrowers who in reality used a broker and those who did not, finding a large degree of overlap between them (see figure C.1.1 and a similar figure in the appendix of Iscenko and Nieboer (2018)).

¹⁰This assumption does not allow us to study the consequences of *distorted* financial advice as in Guiso, Pozzi, Tsoy, Gambacorta, and Mistrulli (2017), so we can interpret our estimates as the upper bound on the value of brokers – if in addition to charging commissions they also provided suboptimal advice, their impact on consumer welfare would be lower. In other words, we abstract from the fact that brokers can be facing potential conflicts of interest between providing the best advice and being compensated by the lender, as discussed by Inderst and Ottaviani (2012c) and Woodward and Hall (2012).

¹¹Early theoretical models which consider price dispersion in markets where some consumers can access sellers directly or use such a clearinghouse include Salop and Stiglitz (1977), Rosenthal (1980), Varian (1980), and Baye and Morgan (2001) among others.

¹²We allow the support to be different for different \mathbf{x}^H .

Since we are considering a market with posted prices, \mathbf{x}^H is a vector of covariates which includes key characteristics of the mortgage¹³, but could also include some elements of \mathbf{x}^G if price discrimination or bargaining are an important feature of the market (Allen et al., 2017). While direct price negotiation is not a typical feature of the UK mortgage market¹⁴, its effects are somewhat mimicked by the fact that lenders typically have broad product menus¹⁵ and it is virtually costless to introduce a new mortgage with a slightly different rate/fees combination. Therefore, the fact that the marginal cost is transaction-specific (i.e. varies across lenders and borrowers) should be seen as an approximation to residual product differentiation which is not captured by the conditioning variables.

3.3.1 Borrowers

Having drawn their search cost, borrowers decide whether to engage in a non-sequential search or use a broker. The search technology is such that i chooses the optimal number of price draws, k , to solve:

$$\min_{k \geq 1} (k - 1)\kappa_i + \mathbb{E} [p_{(1:k)} | \mathbf{x}^G, \mathbf{x}^H] \quad [3.3]$$

Just like in Hong and Shum (2006), we assume that the first draw is costless¹⁶, and the valuation of all consumers is equal to the upper bound of the support of the marginal costs (highest observed price). $\mathbb{E} [p_{(1:k)} | \mathbf{x}^G, \mathbf{x}^H]$ is the expected lowest among k prices drawn from the equilibrium distribution $\mathcal{F}(p | \mathbf{x}^G, \mathbf{x}^H)$, which arises as a result of lenders' profit-maximising pricing decisions given borrowers optimal search behaviour. Unlike Burdett and Judd (1983) where firms and consumers are *ex ante* identical, the equilibrium price dispersion arises both as a result of search and lender heterogeneity.

The cost of using an intermediary is the expected rate paid for the mortgage suggested by the broker plus any commission charged for using the service:

$$\mathbb{E} [p^B | \mathbf{x}^G, \mathbf{x}^H] + \varrho(\mathbf{x}^G, \mathbf{x}^H) \quad [3.4]$$

¹³Specifically: the mortgage term, LTV band, FTB status, duration of the initial deal, indicators whether it is a flexible or cashback mortgage.

¹⁴See the discussion in Benetton (2018) and for anecdotal evidence that some rates are negotiated check <https://www.theguardian.com/money/2013/nov/19/secret-remortgage-rates-special-customers> (accessed 15/09/2018).

¹⁵For example, Kashyap and Rostom (2018) report that the median borrower can choose from 19 different loans within the same lender.

¹⁶In the context of our application, this could be interpreted as the offer from the bank that the consumer has a current account with. In that sense, this brings our model closer to Allen et al. (2017), where borrowers first receive a free offer from their home bank and then decide whether to search or not. The main difference is that we do not assume that the home bank has a cost advantage over the competitors.

Under the assumption that brokers inform the borrower of the best possible deal, we can treat them as auctioneers holding reverse first-price auctions. Therefore, $\mathbb{E}[p^B | \mathbf{x}^G, \mathbf{x}^H] = \mathbb{E}[p_{(1:J)} | \mathbf{x}^G, \mathbf{x}^H]$, so the price is the expected price obtained by searching all J lenders.

Let $k^*(\kappa_i)$ be the optimal number of searches for an individual with unit search cost equal to κ_i . Then the choice of direct search versus using an intermediary is the solution to the following cost minimisation problem:

$$\min_{\{\text{Broker, Direct}\}} \left\{ \mathbb{E}[p^B | \mathbf{x}^G, \mathbf{x}^H] + \varrho(\mathbf{x}^G, \mathbf{x}^H), (k^*(\kappa_i) - 1)\kappa_i + \mathbb{E}[p_{(1:k^*)} | \mathbf{x}^G, \mathbf{x}^H] \right\}, \quad [3.5]$$

Following the insight of Hong and Shum (2006) and lemma 1 in Salz (2017), due to the linearity of search costs in the number of searches and the fact that $\mathbb{E}[p_{(1:k)} | \mathbf{x}^G, \mathbf{x}^H]$ is non-increasing in k , the borrowers in equilibrium will endogenously sort themselves into types defined by the number of searches by forming cut-off points along the search cost distribution:

$$0 \leq \kappa_J(\mathbf{x}^G, \mathbf{x}^H) < \kappa_{J-1}(\mathbf{x}^G, \mathbf{x}^H) < \dots < \kappa_k(\mathbf{x}^G, \mathbf{x}^H) < \kappa_{k-1}(\mathbf{x}^G, \mathbf{x}^H) < \bar{\kappa}(\mathbf{x}^G, \mathbf{x}^H) \leq \infty$$


In the above, κ_k should be understood as the highest search cost so that everyone with $\kappa \in [\kappa_k, \kappa_{k-1}]$ searches exactly k firms. Because we are assuming that brokers are used by individuals with high search cost, $\bar{\kappa}$ is the search cost of the consumer who is indifferent between searching $k - 1$ firms and delegating her search efforts to a broker.

Define $\Delta^B(\mathbf{x}^G, \mathbf{x}^H) = 1 - \mathcal{G}(\bar{\kappa}(\mathbf{x}^G, \mathbf{x}^H) | \mathbf{x}^G)$ as the proportion of borrowers who access the mortgage using a broker. Let $\mathbf{q} = [q_1(\mathbf{x}^G, \mathbf{x}^H), \dots, q_J(\mathbf{x}^G, \mathbf{x}^H)]$ be the proportions of borrowers who search $1, \dots, J$ times. Figure 3.1 below illustrates the equilibrium sorting of borrowers according to the number of searches.

3.3.2 Lenders and equilibrium

Assume that all $j = 1, \dots, J$ lenders have an equal probability of being sampled in the search market. In the price-setting model we depart from two key assumptions used by Salz (2017). Firstly, firms are not able to set different prices in the search and brokered markets, which is reminiscent of the notion of price coherence of Edelman and Wright (2015). Secondly, the same firms participate in the direct and brokered markets. Since we assumed that non-price characteristics were chosen by the borrowers outside of the model, the total profit function is additively separable in profits from every single transaction. Therefore, the same pricing game is played in each of the product-markets defined by the conditioning variables, so for clarity of exposition, we suppress

Figure 3.1: Equilibrium sorting.



Note: Figure shows equilibrium sorting of buyers into types defined by the number of searches and use of brokers. The proportion of borrowers who use brokers is the blue shaded area under the search cost density. The areas between the cutoffs on the x-axis determine the proportions of buyers who search different number of firms. PDF not drawn up to scale.

the conditioning sets. Firm j with marginal cost c_{ij} solves:

$$\max_p \Delta^B \cdot \Pi^B(p, c_{ij}) + (1 - \Delta^B) \cdot \Pi^D(p, c_{ij}; \mathbf{q}), \quad [3.6]$$

where $\Pi^B(\cdot)$ and $\Pi^D(\cdot)$ are the profits in the brokered and direct search market, respectively. Suppose all lenders have equal probabilities of being found by borrowers. Then the probability of being searched by a borrower who samples ℓ lenders is $\frac{\ell}{J}$ and we can restrict our attention to symmetric equilibria. Following Burdett and Judd (1983) and assuming that brokers hold first-price auctions, we can rephrase the problem as:

$$\max_p \Delta^B \cdot (p - c_{ij})(1 - \mathcal{F}(p))^{J-1} + (1 - \Delta^B) \cdot (p - c_{ij}) \sum_{\ell=1}^J q \frac{\ell}{J} (1 - \mathcal{F}(p))^{\ell-1} \quad [3.7]$$

Due to costs being lenders' private information drawn from the same distribution, the probability that a lender with $\ell - 1$ competitors wins the contract (is the cheapest amongst ℓ firms) is $(1 - \mathcal{F}(p))^{\ell-1}$.

Let $\tilde{q}_\ell = (1 - \Delta^B)q_\ell$ for $\ell = 1, \dots, J-1$, and $\tilde{q}_J = \Delta^B + (1 - \Delta^B)q_J$. Then the maximisation problem simplifies to:

$$\max_p (p - c_{ij}) \sum_{\ell=1}^J \tilde{q}_\ell \frac{\ell}{J} (1 - \mathcal{F}(p))^{\ell-1}, \quad [3.8]$$

which is the same as the price-setting problem in the pure search problem considered in chapter 2 with distorted search probabilities $\tilde{\mathbf{q}} = [\tilde{q}_1, \dots, \tilde{q}_J]$. Let $\beta(c_{ij}; \tilde{\mathbf{q}})$ denote the optimal strategy given beliefs about borrowers' search decisions. Using the envelope theorem with the boundary condition that $\beta(\bar{c}; \tilde{\mathbf{q}}) = \bar{c}$ yields the optimal pricing strategy:

$$\beta(c_{ij}; \tilde{\mathbf{q}}) = c_{ij} + \frac{\sum_{\ell=1}^J \tilde{q}_\ell \ell \int_{s=c_{ij}}^{\bar{c}} (1 - \mathcal{H}(s))^{\ell-1} ds}{\sum_{\ell=1}^J \tilde{q}_\ell \ell (1 - \mathcal{H}(c_{ij}))^{\ell-1}}. \quad [3.9]$$

The equilibrium price distribution, \mathcal{F} , emerges as a result of lenders pricing according to (3.9).

We are now ready to define a symmetric Bayesian-Nash equilibrium of the game:

DEFINITION . *The pair $(\tilde{\mathbf{q}}, \beta(\cdot; \tilde{\mathbf{q}}))$ is a symmetric Bayesian Nash equilibrium if:*

- (i) *for every $\tilde{\mathbf{q}}$ when all firms apart from j use pricing strategy $\beta(\cdot; \tilde{\mathbf{q}})$, $\beta(\cdot; \tilde{\mathbf{q}})$ is the best response for firm j ;*
- (ii) *given the price distribution induced by $\beta(\cdot; \tilde{\mathbf{q}})$, $\tilde{\mathbf{q}}$ is a vector of proportions of consumers' optimal search.*

We restrict our attention to monotone pure-strategy equilibria. The action space is compact and the pricing functions are strictly increasing in cost so the existence results from Reny (2011) apply. In general, lenders' payoff function can be seen as a mixture of two auctions – one where all firms participate (broker), and one where the number of competitors is unknown (direct search). Mixing probabilities are then determined in equilibrium by optimal search decisions made by borrowers.

3.4 Identification and estimation

This section discusses the identification of the model's primitives, that is the set of conditional search cost distributions $\mathcal{G}(\cdot | \mathbf{x}^G)$ and distribution of cost of providing the loan, $\mathcal{H}(\cdot | \mathbf{x}^H)$. We argue that the model imposes enough structure on the data for the aforementioned distributions to be nonparametrically identified.

Throughout the section we assume that for each mortgage, we observe the price, p_{ij} , whether

it was accessed directly or via a broker¹⁷, and the values of the borrower and loan characteristics, \mathbf{x}^G and \mathbf{x}^H , respectively. In addition to that, we assume that we can observe (or construct from the data) lender market shares for each of the combinations of conditioning variables and broker commissions. The goal of nonparametric identification is then to establish a mapping from these data to the unobserved primitives using the theoretical restrictions imposed by the model. The main identification theorems for a pure search model are presented in chapter 2 and appendix B, and draw on findings from the literature on nonparametric auction estimation, in particular Guerre, Perrigne, and Vuong (2000).¹⁸ We therefore only devote space in this section to emphasise certain aspects of the identification strategy which have not been discussed in previous literature.

3.4.1 Differences from Salz (2017)

The identification strategy and estimation techniques used here differ from those employed by Salz (2017) due to altered assumptions on the supply-side of the model. This section discusses the differences.

The main point of departure is that we do not assume that the composition of firms offering their products directly and through brokers is different. In our setting, the same lenders offer their loans directly or through a broker. Because we only focus on the biggest six lenders, we do not rely on the *ex ante* classification of firms into high and low types, which is one of the fundamental identifying assumptions made by Salz (2017).¹⁹ We therefore restrict our attention to symmetric equilibria, where all firms have the same underlying cost distribution. This allows us to use the data from both brokered and direct mortgages when estimating \mathcal{H} , unlike Salz (2017) who used only brokered data to recover firms' cost distributions.

The second difference is that we do not work with residualised prices²⁰, so that we do not need to make the assumption that the cost is additively separable in a linear index of characteristics. Instead, we nonparametrically estimate cost separately for each combination of \mathbf{x}^H . This allows to capture potential nonlinear patterns pricing and underlying cost distributions. For example, with the linear index restriction, the estimated cost distributions for 70% LTV and 90% LTV loans could differ only in their means while the fully nonparametric approach we adopt here allows also for different higher moments.

¹⁷In practice the information on the proportion of brokered loans should suffice.

¹⁸See also Athey and Haile (2007) for a comprehensive overview.

¹⁹The identifying assumption in Salz (2017) is that the econometrician needs to observe at least one firm of type H and one of type L in the brokered and direct markets. Then one can use the structure imposed by broker auctions to recover the distribution of costs for low and high type firms.

²⁰That is residuals from a hedonic regression of prices on a vector of loan characteristics.

Finally, to identify the distribution of search costs, our approach uses data on prices and market shares and a technique that minimizes the distance between market shares predicted by the model and the data. We then use the estimated proportions of borrowers searching different number of lenders to estimate the distributions of marginal costs. Salz (2017), on the other hand, first estimates the distribution of costs from prices of brokered contracts, and then uses those estimates to recover proportions of businesses (consumers in his model) searching different number of waste disposal firms. Our approaches are therefore non-nested and valid under different sets of assumptions.

3.4.2 Role of exclusion restrictions

In the exposition of the model and the ensuing empirical analysis, we refer to two types of conditioning variables: \mathbf{x}^G shift the consumer search cost distribution and can be thought of as variables defining consumer type. \mathbf{x}^H are primarily mortgage characteristics that directly affect the cost of providing the loan from the lenders' perspective. The two sets of variables can contain some common elements, or, in the most extreme scenario, fully overlap.

A natural question is whether lack of exclusion restrictions precludes identification of the unobserved cost distributions. In general, the answer is no, but it might lead to a situation where the distribution of search cost is only identified at very few points on its support. This discussion is related to the finding in Hong and Shum (2006) which was further elaborated by Moraga-González, Sándor, and Wildenbeest (2013) – namely, absent any other dimension of variation in the data (e.g. across local markets or time), one is unable to identify the search cost distribution beyond a set of $J - 1$ points where J is the number of firms. A solution to this problem is to pool estimates from multiple markets, as shown e.g. by Sanches et al. (2016a).

Exclusion restrictions help generate such markets. In principle for each consumer type (described by \mathbf{x}^G) we can pool estimates from different mortgage types (\mathbf{x}^H). To provide a specific example²¹, assuming that all first-time buyers, aged 30+, who live in cities and their income is above the median have the same distribution of search cost, one can first obtain different sets of estimates for each type of mortgage and then combine them to obtain a smoothed version of the search cost CDF. Absent any other variation in the data, one might have to resort to a parametric specification to be able to conduct meaningful counterfactual inference.

²¹Moraga-González et al. (2013) suggest that pooling data is possible across much more heterogeneous markets than the ones in our application: "(...) to estimate the costs of search in the market for carpentry, one could pool data from the various professional services needed to refurbish a house: a carpenter, an electrician, a painter, a plumber, a bricklayer, a tiler, etc."

3.4.3 Estimation steps

We now describe the main steps of the estimation algorithm based on the constructive identification presented in chapter 2. In short, we begin with obtaining the empirical distributions of prices, which are then used in conjunction with data on market shares to back out the proportion of borrowers searching different number of lenders. In this step, we need to impose the constraint that the proportion of borrowers who know offers from all lenders is no smaller than the proportion of brokered loans in the data. Next, as discussed in the preceding subsection, we pool the data across markets to obtain the estimate of the full search cost CDF. Finally, we use equilibrium bidding function to construct pseudo-costs and then use kernel techniques to obtain their density.

To ease the notation, let s index distinct discrete combinations of $(\mathbf{x}^G, \mathbf{x}^H)$ and, for the sake of brevity, let $\mathcal{F}_s(p) \equiv \mathcal{F}(p|s)$, $f_s(p) \equiv f(p|s)$, and $\Delta_s^B \equiv \Delta^B(s)$. The estimation algorithm consists of the following 5 steps:

1. In the first step we estimate $\mathcal{F}_s(p)$ and $f_s(p)$, separately for all $s \in \{\mathbf{x}^G \times \mathbf{x}^H\}$, where the cardinality of the set of covariates depends on how the variables are discretized and can potentially be very large. The CDF is estimated simply as:

$$\hat{\mathcal{F}}_s(p) = \frac{1}{n_s} \sum_{i=1}^{n_s} \mathbf{1}\{p_i \leq p\}.$$

To estimate the density, we need to address the problem of bias near the lower boundary and the possibility that the density near the upper boundary may be unbounded. To tackle this, we use an asymmetric Beta kernel suggested by Chen (1999) that performs well on densities defined over compact supports²² together with the transformation method of Marron and Ruppert (1994) near the upper boundary.

2. Using our transaction data, we construct a vector of observed market shares²³, $Y_s = (Y_{1s}, \dots, Y_{J_s})^\top$ for every s . Let \mathbf{X}_s be a $J_s \times J_s$ matrix such that $(\mathbf{X}_s)_{j\ell} = \frac{\ell}{J_s} (1 - \mathcal{F}_s(p_j))^{\ell-1}$. Then under Assumption I from chapter 2 we have:

$$\tilde{\mathbf{q}}(s) = \frac{\mathbb{E}[\mathbf{X}_s^\top \mathbf{X}_s]^{-1} \mathbb{E}[\mathbf{X}_s^\top Y_s]}{\iota^\top \mathbb{E}[\mathbf{X}_s^\top \mathbf{X}_s]^{-1} \mathbb{E}[\mathbf{X}_s^\top Y_s]}, \quad [3.10]$$

where ι denotes a vector of ones. To obtain an estimate of $\tilde{\mathbf{q}}_s$, we use the estimated CDFs from step 1 and evaluated them at the average price charged by firm j conditional on s ,

²²The implementation comes from the `npuniden.boundary` function from the `np` package in R (Hayfield and Racine, 2008).

²³We experimented with both aggregate market shares over the entire period of the sample as well as quarterly shares. With a fine grid for $(\mathbf{x}^G, \mathbf{x}^H)$, obtaining precise estimates of quarterly shares requires a lot of data to prevent

obtaining $(\widehat{\mathbf{X}}_s)_{j\ell} = \frac{\ell}{J_s} \left(1 - \widehat{\mathcal{F}}_s(\bar{p}_j)\right)^{\ell-1}$ as the sample analogue of the \mathbf{X}_s matrix.²⁴ The intuition behind this step is that, conditional on price, the observed difference in market shares can only be explained by some consumers having different number of offers to compare than others. Therefore the variation in market shares conditional on price identifies the search proportions.

To accommodate the restriction that $\tilde{q}_{J_s}(s) \geq \Delta_s^B$, we use constrained quadratic programming to solve the least squares problem (3.10). The right-hand side of the constraint, Δ_s^B is the proportion of brokered mortgages and can be directly obtained from the data.

3. Estimate the vectors of cutoff types $\boldsymbol{\kappa}_s \equiv \boldsymbol{\kappa}(s)$ for each $s \in \{\mathbf{x}^G \times \mathbf{x}^H\}$, where for $\ell \in \{1, \dots, J_s - 1\}$:

$$\kappa_{\ell}(s) = \mathbb{E}_{\mathcal{F}_s} [p_{(1:\ell)}] - \mathbb{E}_{\mathcal{F}_s} [p_{(1:\ell+1)}]$$

and the marginal type who is indifferent between using a broker and searching directly is estimated as:

$$\bar{\kappa}(s) = \frac{\varrho(s) - \left(\mathbb{E}_{\mathcal{F}_s} [p_{(1:k^*)}] - \mathbb{E}_{\mathcal{F}_s} [p_{(1:J)}]\right)}{k^* - 1}.$$

$\varrho(s)$ is the average broker commission and k^* is the equilibrium number of searches of the marginal type. To determine k^* , we find the lowest ℓ , such that $\kappa_{\ell}(s) < \bar{\kappa}(s)$.²⁵ To estimate the expectations of the order statistics, we draw repeatedly from the price distributions and calculate the sample averages of the minimum prices.

Finally, $\hat{\mathbf{q}}(s)$ estimated in the previous step can be used to recover \mathcal{G} evaluated at the cutoff points as follows:

$$\begin{aligned} \mathcal{G}(\bar{\kappa}(s)|\mathbf{x}^G) &= 1 - \Delta_s^B \\ \mathcal{G}(\kappa_{k^*}(s)|\mathbf{x}^G) &= 1 - \Delta_s^B - \hat{q}_{k^*}(s) \\ &\vdots \\ \mathcal{G}(\kappa_{J_s-1}(s)|\mathbf{x}^G) &= 1 - \Delta_s^B - \hat{q}_{k^*}(s) - \dots - \hat{q}_{J_s-1}(s) \end{aligned}$$

²⁴The identifying assumption suggested in chapter 2 is that the observed market shares are systematically related (proportional) to the ex-ante probabilities of winning the procurement auction. A slight difficulty in the empirical application using transaction data is that constructing market shares from transaction data typically requires summing over multiple transactions by the same firm, which tend to be associated with different prices. Therefore one needs to choose at which price should the CDF be evaluated. Using the average is consistent with the proportionality assumption – since lenders are assumed to have the same underlying cost distribution, we can only explain differences in aggregate market shares in the data by lower/higher draws from \mathcal{H} and consequently lower/higher average prices quoted.

²⁵Clearly, $\bar{\kappa}(s)$ is not identified if $k^* = 1$, so when someone who now is indifferent between using a broker or not would not search beyond the first offer she receives for free if intermediation was not available. In such a case, we replace $\bar{\kappa}(s) = \kappa_1(s) + \epsilon$ where $\epsilon \sim \text{Unif}[0, \kappa_1(s)]$.

4. Let $\{s^H(x)\}_{x=1}^{|\mathbf{x}^G|}$ denote the collection of partitions of $\{\mathbf{x}^G \times \mathbf{x}^H\}$ such that $\mathbf{x}^G = x$. The cardinality of each of those sets, $|s^H(x)|$, is equal to the cardinality of \mathbf{x}^H , say n_H . For each of the sets we now have n_H estimates of the cutoffs $\{\kappa^t\}_{t \in s^H(x)}$ and $\{\mathcal{G}(\kappa^t | \mathbf{x}^G)\}_{t \in s^H(x)}$. We then pool the estimates using the method suggested in section 4 of (Sanches et al., 2016a). Specifically, separately for each x , we seek to minimize the following least squares criterion function:

$$\Psi_x(g) = \frac{1}{n_H} \sum_{t=1}^{n_H} \sum_{\ell=1}^{J_t-1} [\mathcal{G}(\kappa_\ell^t | \mathbf{x}^G = x) - g(\kappa_\ell^t)]^2,$$

where g is a flexible function of the cutoffs. To impose appropriate shape restrictions on the estimated CDF, we choose Bernstein polynomials²⁶ to construct the sieve.

This step results in a sieve-least squares estimator for $\mathcal{G}(\cdot | \mathbf{x}^G)$, whose theoretical properties and assumptions needed for consistency are discussed in Sanches et al. (2016a).

5. In the final step we recover the distributions of lenders' marginal costs. This step is reminiscent of recovering the distribution of valuations from observed bids in a first-price auction (Guerre et al., 2000). First, for each observed price, we construct pseudo-marginal costs using the inverse of the bidding function:

$$\hat{c}_{ij}(s) = p_{ij} - \frac{\sum_{\ell=1}^{J_s} \hat{q}_\ell \ell \left(1 - \hat{\mathcal{F}}_s(p_{ij})\right)^{\ell-1}}{\hat{f}_s(p_{ij}) \sum_{\ell=1}^{J_s} \hat{q}_\ell \ell (\ell - 1) \left(1 - \hat{\mathcal{F}}_s(p_{ij})\right)^{\ell-2}}, \quad [3.11]$$

As before, let $\{s^G(z)\}_{z=1}^{|\mathbf{x}^H|}$ be defined as the collection of partitions of $\{\mathbf{x}^G \times \mathbf{x}^H\}$ such that $\mathbf{x}^H = z$. We can now pool the generated pseudo-costs corresponding to each value of \mathbf{x}^H : $\{\hat{c}_{ij}(t)\}_{t \in s^G(z)}$ and proceed to estimate $\mathcal{H}(\cdot | \mathbf{x}^H)$ and $h(\cdot | \mathbf{x}^H)$. As with the price density, we estimate the density using boundary kernels to reduce the bias.

3.5 Results

We now provide an overview of the estimation results obtained by applying our structural model and estimation strategy to the PSD data. Since we are interested in recovering search and marginal cost distributions conditional on observed heterogeneity, the number of different search cost distributions is equal to the number of bins defined by the chosen partition of \mathbf{x}^G , and, correspondingly,

²⁶A Bernstein polynomial of order P is a set of $p = 0, \dots, P + 1$ functions where $g_{pP}(\kappa) = \frac{P!}{p!(P-p)!} \kappa^p (1 - \kappa)^{P-p}$

the number of distinct marginal cost distributions is equal to the cardinality of \mathbf{x}^H .

As discussed in section 3.4.2, exclusion restrictions are not necessary for theoretical identification of the model's primitives,²⁷ they turn out to be useful in practice. With some variables excluded from \mathbf{x}^G but included in \mathbf{x}^H , we are able to pool search cost cutoff estimates and the corresponding values of the survival function originating from different *markets* defined by different \mathbf{x}^H . This allows us to identify the distribution on a wider support, instead of only on a small collection of discrete points.²⁸

Borrower and loan characteristics used in the structural model are displayed in table 3.4. Their choice was primarily driven by the tradeoff between computational feasibility and willingness to accommodate rich borrower and product-level heterogeneity. Since the model needs to be solved for each combination of $(\mathbf{x}^G, \mathbf{x}^H)$, we rely on a rather coarse discretization of continuous variables. A related issue is that for kernel methods to provide a reliable estimate of the pdf of observed prices we need possibly many data points in each of the bins. Therefore, out of the initial 27,648 bins we used only those with 50 or more observations. This leaves us with 3,697 combinations (86.68% of the total number of mortgages in our main sample) representing the most popular products and borrower types. While by doing this we are no longer working with the entire universe of mortgages, the scope of loans we look at is still much broader than in previous literature using search models to study mortgage markets.²⁹

Table 3.4: Covariate selection.

Variable	Discretization	# bins
\mathbf{x}^G (16 combinations)		
Age	<30, 30+	2
Income	Below median, Above median	2
FTB status	FTB, Non-FTB	2
Location	Urban, Rural	2
\mathbf{x}^H (1,728 combinations)		
LTV	≤70, 71-75, 76-80, 81-85, 86-90, 91-95	6
Deal length	2-, 3-, 5-year	3
Duration	<10, (10;15], (15;20], (20;25], (25;30], (30;35]	6
Loan value	4 quantiles	4
Flexible	Yes, No	2
Cashback	Yes, No	2
Total: 27,648 bins		

Note: Table presents the selection of conditioning variables used in the estimation of the structural models. The total number of bins is the cardinality of the Cartesian product of the elements of \mathbf{x}^G and \mathbf{x}^H .

We use the definition of monthly mortgage cost from 3.2.3 as the price variable. To remove

²⁷ \mathcal{G} and \mathcal{H} are identified even if $\mathbf{x}^G = \mathbf{x}^H$.

²⁸See Moraga-González et al. (2013) for a discussion on identifying search cost distributions using data from multiple markets and Sanches et al. (2016a) for the theoretical properties of the method employed here.

²⁹For example, Allen, Clark, and Houde (2013) look exclusively at FTBs taking out loans with 25 year amortisation and 5-year initial deal period.

any dispersion stemming from macroeconomic shocks (e.g. changes to the BoE interest rates), we detrend the prices and denominate them in January 2016 GBP by taking the residual from regressing prices separately within each (x^G, x^H) cell on a full set of monthly dummies. Ultimately, the monetary magnitudes of all results in this and the following section should be interpreted relative to a monthly interest and fee payment on a median-sized mortgage denominated in prices from the beginning of 2016.

3.5.1 Borrowers' search costs

Our main estimation results are summarised in table 3.5. The estimated quantities represent the monthly unit cost of contacting an additional bank and obtaining a price quote. The median costs range from £22.04 (young first time buyers living in urban areas) to £107.91 (low income, non-first time buyers aged 30+ in urban areas). In relative terms, they represent between 5.85 and 35.1% of the median interest-only payment.

Table 3.5: Summary statistics for estimated search cost distributions.

#	x^G bin	Median	Median searchers	Med(κ)/ \bar{p}
1	Age <30 — Low income — FTB — Rural	23.38	16.68	5.87%
2	Age 30+ — Low income — FTB — Rural	56.11	20.32	14.55%
3	Age <30 — High income — FTB — Rural	47.07	16.78	13.01%
4	Age 30+ — High income — FTB — Rural	23.04	14.97	6.63%
5	Age <30 — Low income — Non-FTB — Rural	42.38	22.94	12.14%
6	Age 30+ — Low income — Non-FTB — Rural	55.07	21.60	16.71%
7	Age <30 — High income — Non-FTB — Rural	45.64	12.96	14.30%
8	Age 30+ — High income — Non-FTB — Rural	46.32	22.51	14.91%
9	Age <30 — Low income — FTB — Urban	28.77	8.73	7.01%
10	Age 30+ — Low income — FTB — Urban	50.95	10.80	12.74%
11	Age <30 — High income — FTB — Urban	22.04	6.39	5.85%
12	Age 30+ — High income — FTB — Urban	24.03	6.12	6.55%
13	Age <30 — Low income — Non-FTB — Urban	101.77	18.52	28.59%
14	Age 30+ — Low income — Non-FTB — Urban	107.91	18.02	32.17%
15	Age <30 — High income — Non-FTB — Urban	31.87	14.95	9.76%
16	Age 30+ — High income — Non-FTB — Urban	104.80	17.80	33.82%

Note: Table presents selected features of nonparametrically estimated search cost distributions for 16 different borrower types (referred to as x^G bins). The third column contains the median search cost in GBP/month in the initial period. Fourth column shows the median search cost among borrowers who do not use brokers. Fifth column shows the ratio of the median search cost to the median and mean price paid.


These results are higher, but still quantitatively similar to other estimates provided in the literature. For instance, Allen et al. (2017) estimate the mean search cost to be \$29/month, while Agarwal et al. (2017) convert their estimate expressed in basis points to \$27/month on a representative loan. The main reason why our estimates are higher is because of the assumption that consumers with high search cost use brokers and the fact that on average over 70% of the mortgages in our data are brokered. This is also the reason why higher percentiles of the search cost distribution are not identified³⁰ – the data are informative about the fraction of borrowers with

³⁰This is why in some cases we could not calculate the interquartile range.

$\kappa > \bar{\kappa}$ but without parametric assumptions we cannot identify the shape of the distribution above $\bar{\kappa}$. What helps to some extent is variation in $\bar{\kappa}$ induced by different combinations of \mathbf{x}^H , yet this does not allow for identification of \mathcal{G} on $[0, \infty)$.

The median search costs among the borrowers who end up are substantially lower, ranging from £6 to about £22.51. The part of the distributions corresponding to searchers are also plotted in figure 3.2. The results point to a high degree of heterogeneity across different demographics. We also find that some of the distributions are clearly bimodal, with the first peak below £10, showing that the consumers who do not use brokers can efficiently search on their own. To compare the distributions across different demographics, we constructed an additional array of graphs in which we compare distributions across one trait keeping the other ones fixed (see figure C.2.2 in appendix B). The two main findings that emerge from that comparison are that borrowers from rural areas seem to have higher search costs than inhabitants of cities and, perhaps surprisingly, FTB’s search costs are stochastically dominated by the costs of non-FTBs.

Figure 3.2: Estimated search cost CDFs.



Note: Search cost denominated in January 2016 GBP (£) per month. Bernstein sieves were used to impose shape restrictions (non-decreasingness). The respective distributions are identified on $[\kappa_{min}(\mathbf{x}^G, \mathbf{x}^H), \kappa_{max}(\mathbf{x}^G, \mathbf{x}^H)]$, that is the lowest and highest cutoff estimated in the data.

The first finding is intuitive – since search involves interviews at bank branches, the process is more costly for borrowers living in areas with low presence of lenders. The second observation is trickier to interpret, as it rejects the possible role of learning in the reduction of future search

costs. However, since the interval between the first and subsequent purchase or remortgage is typically several or more years, experience from the first purchase might not be playing a major role. Moreover, given that our data set does not include borrowers refinancing with their previous loan provider, our estimate of the search cost might also be absorbing part of the otherwise unobserved switching cost.³¹

To our surprise, we did not find any unambiguous effects of age or income, which can be seen as proxies for financial literacy (Hastings et al., 2013). One potential explanation can be that our definition of search cost does not separate between cognitive and opportunity cost and the two effects might be counterbalancing each other. To assure that our findings are robust to different ways of defining age and income bins, we also estimated search costs with more finely discretized grids. Namely, we considered 4 different age buckets and 4 income quartiles, resulting in 64 different distributions. This did not really provide us with any additional insights into the effects of age and income on search cost, while keeping the conclusions on the effect of FTB status and location virtually unchanged. We therefore decided to stick with the more parsimonious specification.

3.5.2 Lenders' costs and margins

We now present the estimates of the supply-side primitives, i.e. lenders' marginal costs and associated markups. Instead of showing estimated CDFs for all possible combinations of conditioning variables, \mathbf{x}^H , we only look at some aggregate statistics of marginal distributions to make sure that the estimates align with common sense. For example, we expect mortgages that are riskier from the banks' perspective (e.g. higher LTV, longer term) to be associated with higher risk premia which are expected to be captured by the cost estimates. Table 3.6 summarises our main results.

Intuitively, the averages of the estimated costs increase in LTV and length of the amortisation period, and are higher for loans with 5-year fixed rate period than for shorter periods of fixed rate. Given that in our definition of price we normalised the quantity of the loan to correspond to the median value, it is not surprising that we do not find major differences in costs across the four quartiles of the loan value distribution. We do not find major differences in the cost distributions for mortgages offering flexible repayment schemes or cashback. Densities of marginal³² distributions of costs are shown in figure 3.3. Clearly, LTV level as the main indicator of loan's riskiness is the main driver of cost. The third figure in the top panel exhibits an increase in cost variance as

³¹Several authors studied the role of switching costs in the banking industry – see e.g. Kim et al. (2003), Deuffhard (2016), Honka et al. (2017).

³²Here we use *marginal* in the context of a statistical definition of a marginal distribution.

Table 3.6: Summary statistics for estimated marginal costs and margins.

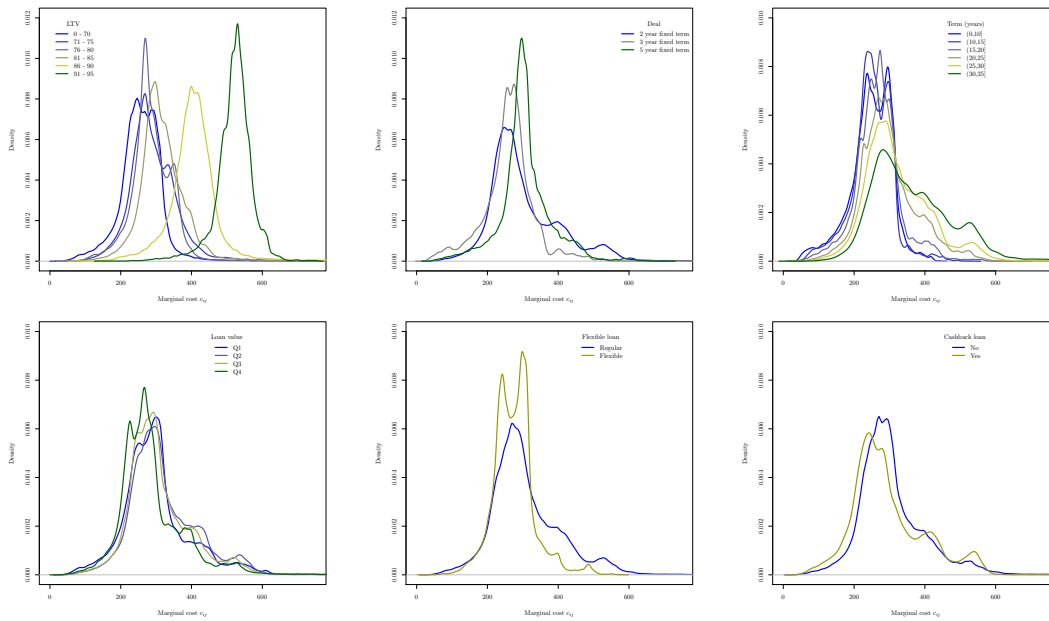
x^H category	Marginal cost			Price-cost margin		
	Mean	Median	IQR	Mean	Median	IQR
LTV						
≤ 70	258.9	258.2	66.8	14.07	8.45	11.26
71-75	293.3	282.8	80.4	9.89	6.22	7.17
76-80	288.3	281.8	66.2	11.09	6.82	8.26
81-85	318.8	309.2	69.0	9.04	5.65	6.96
86-90	407.8	405.1	64.6	8.75	5.65	5.79
91-95	519.9	524.8	52.5	5.12	2.95	3.32
Deal						
2 years	305.6	276.6	117.2	10.60	6.58	8.07
3 years	260.5	263.0	66.4	19.61	12.90	15.18
5 years	306.3	300.7	62.0	13.44	7.63	10.6
Term						
≤ 10 years	248.7	253.6	74.4	17.98	11.38	17.63
(10;15]	251.2	252.2	66.9	16.63	10.20	15.05
(15;20]	267.7	266.8	68.7	14.16	8.67	11.56
(20;25]	298.0	285.6	91.3	11.22	6.96	7.68
(25;30]	323.1	303.4	116.3	9.62	6.09	6.54
(30;35]	369.1	343.0	155.6	7.16	4.54	5.59
Value						
Q1	302.7	288.7	91.8	15.07	9.39	13.21
Q2	321.1	301.0	113.4	11.23	6.98	8.24
Q3	309.8	291.5	98.9	10.13	6.05	7.36
Q4	287.6	269.1	92.4	10.08	6.13	6.93
Flexible						
Regular	309.5	289.0	109.8	12.04	7.15	9.25
Flexible	275.5	274.0	66.7	9.30	5.89	7.24
Cashback						
No cashback	306.3	288.5	97.1	11.14	6.63	8.24
Cashback	294.5	272.8	118.8	14.85	9.97	12.43

Note: Means, medians and interquartile ranges of estimated marginal cost and price-cost margins defined as $PCM_{ij} = \frac{p_{ij} - c_{ij}}{p_{ij}}$. Costs expressed in £, PCMs in %.

the mortgage duration increases. Since most mortgagors refinance, the loans with 25+ years are mostly those of first-time buyers, so we can also interpret this finding in terms of the cost of servicing loans being more idiosyncratic for first-time borrowers who amortise over longer periods of time.


The distribution of markups is right-skewed with an average of 11.64% and median of 6.97% (see figure 3.4 and more detailed results in figure C.2.1 in appendix B). Despite high market concentration, the six biggest banks do not seem to be able to exert substantial market power. While the estimates might appear to be low, it is worth emphasising that the model does not define the PCM in the canonical way as the markup over the interbank swap rate (such as LIBOR) or the Bank of England base rate. Instead, since our definition of price includes also the upfront fees and we do not model lenders' fixed costs in any other way, we believe that the estimate is closer to the cost of servicing the loan over the fixed-rate period, including the cost of processing the mortgage application, hedging against the risk of default etc. Our estimates of the markups are higher than those obtained by Allen et al. (2017) in a recent study of the Canadian mortgage market, where the average Lerner index is estimated to be 3.2%. This is still more than 3 times lower than the average and 2 times lower than the median of our estimates.

Figure 3.3: Marginal distributions of marginal costs.



Note: Kernel estimates of marginal costs for each conditioning variable in \mathbf{x}^H integrated over the remaining covariates.

Figure 3.4: Distribution of price-cost margins.



Note: Kernel estimate of the density of price-cost margins defined as $PCM_{ij} = \frac{p_{ij} - c_{ij}}{p_{ij}}$.

Overall, the results in this section are consistent with some recent evidence³³ that in the last few years the market became more competitive and lenders are no longer able to enjoy high margins. From the perspective of the structural model's mechanics, low markups emerge as an artefact of high proportion of borrowers using brokers, whose presence, by construction, stimulates competition between lenders. We will return to this discussion in the following section.

³³See, for example, The Guardian (27/04/2017): *Low rates, tight margins: the mortgage market looks worryingly familiar*: <https://www.theguardian.com/business/2017/apr/23/mortgage-market-lower-rates-tight-margins-worryingly-familiar> (accessed 10/09/2018).

3.6 Counterfactuals

We use the estimates from the previous section to derive our main results. In the current market setting, borrowers possess different information sets due to heterogeneous search costs. Those with low search costs are able to efficiently search through products offered by competing lenders, while those with higher search costs either have only one mortgage to choose from or need to use an intermediary.

In the first experiment, we are interested in quantifying the value of information provided by brokers to consumers with high search costs. To do that, we need to simulate counterfactual market outcomes in a scenario where brokers are not present in the marketplace. Since intermediaries originally reduce the monopoly power of lenders, we expect the prices and markups to be higher in the new equilibrium. Borrowers, on the other hand, will no longer have to pay broker commissions, but their total search expense will change.

The second counterfactual concerns a hypothetical market centralisation. Suppose the regulator establishes a platform, where all lenders need to post their prices and borrowers are matched to the best offer. This setting is equivalent to a pure first-price procurement auction with no search involved. We study how such regulation would affect consumer surplus and the prices set by the lenders.

3.6.1 Value of information provided by brokers

In the absence of intermediation, the borrowers' problem is to choose the optimal number of searches given their search cost. Therefore, to find the new equilibrium, we solve the fixed-point problem defined in the space of (new) optimal search proportions denoted as $\hat{\mathbf{q}}$:

$$\hat{q}_\ell = \begin{cases} 1 - \mathcal{G} \left[\int \mathcal{H}(\xi(p, \hat{\mathbf{q}}) | \mathbf{x}^H) (1 - \mathcal{H}(\xi(p, \hat{\mathbf{q}}) | \mathbf{x}^H)) dp | \mathbf{x}^G \right] & \text{for } \ell = 1 \\ \mathcal{G} \left[\int \mathcal{H}(\xi(p, \hat{\mathbf{q}}) | \mathbf{x}^H) (1 - \mathcal{H}(\xi(p, \hat{\mathbf{q}}) | \mathbf{x}^H))^\ell dp | \mathbf{x}^G \right] & \text{for } \ell > 1 \\ -\mathcal{G} \left[\int \mathcal{H}(\xi(p, \hat{\mathbf{q}}) | \mathbf{x}^H) (1 - \mathcal{H}(\xi(p, \hat{\mathbf{q}}) | \mathbf{x}^H))^{\ell+1} dp | \mathbf{x}^G \right] & \end{cases} \quad [3.12]$$

As discussed in chapter 2, Brouwer's theorem guarantees the existence of a fixed point. While uniqueness cannot be proved, we experimented with different starting points finding that the algorithm converges to the same solution in the interior of the simplex.³⁴ The new search proportions then feed into the firms' pricing functions to generate a set of counterfactual conditional price

³⁴In several cases we observed the FP iteration to converge to a degenerate (monopoly) solution with $\hat{q}_1 = 1$. In such cases we tried a different starting points to find a solution in the interior. If this method failed, we took the monopoly outcome to be the only solution to the problem, otherwise we discarded it.

distributions $\dot{\mathcal{F}}(\cdot|\mathbf{x}^G, \mathbf{x}^H)$ which can then be sampled from.

The auction model assumes that for all consumers, the valuation of the mortgage is at least as high as the upper limit of the support of the cost distribution. Therefore, realised consumer surplus for a borrower paying p is:

$$CS = \bar{v} - p - SE$$

SE stands for search expenditure and is equal to $\kappa(k - 1)$ if the borrower with search cost κ accessed the loan directly by contacting k banks, or ϱ if she used a broker and paid commission equal to ϱ .

Without intermediation,

$$\dot{CS} = \bar{v} - \dot{p} - \dot{SE}$$

where \dot{p} is the new price drawn from $\dot{\mathcal{F}}$ and \dot{SE} is the new realised search expenditure which now does not include the possibility of using a broker. We now define the value of information (VoI)³⁵ as the difference between the expected CS and \dot{CS} :

$$\text{VoI} = \mathbb{E}(CS - \dot{CS}) = \mathbb{E}(\dot{p} - p) + \mathbb{E}(\dot{SE} - SE) \quad [3.13]$$

Tables 3.7 and 3.8 report our main counterfactual results. We estimate the average value of information provided by brokers in this market to be £112.15. Given our definition of price, this means that the existence of brokers helps an average mortgagor save over £112.15 per month in *sunk* expenditures (i.e. those not related to paying off the principal). If brokers were not present in the market, borrowers would be spending on average 33.7% more in prices and 16.33% in search cost. This results suggest that the role of borrowers in limiting lenders' monopoly power which arises when consumers do not search enough is substantial. Importantly, intermediation generates a positive externality for borrowers who search directly.³⁶ We now look at the same figures disaggregated by consumer and mortgage type to see whether everyone is indeed better off.

Table 3.7 breaks down the main result by borrower characteristics. Clearly, young, low-income, first-time buyers benefit the most. The counterfactual price increases they would be facing in a world with no intermediaries reach up to 55.75%. They are also matched by significant

³⁵The definition adopted here is slightly different from e.g. Baye, Morgan, and Scholten's (2006) discussion of the Varian (1980) model who define value of information as the difference between the expected price of consumers who access the clearinghouse and those that do not.

³⁶This is a natural consequence of the price coherence assumption and is somewhat different from the same finding in Salz (2017) who allowed separate price setting in the two market segments.

Table 3.7: Value of information: breakdown by borrower types.

	VOI	$\% \Delta p$	$\% \Delta SE$
Overall	112.15	+33.70%	+16.33%
Age			
<30	199.98	+55.75%	+65.49%
30+	93.16	+28.94%	+5.70%
Income			
Low	125.32	+36.80%	-18.62%
High	99.31	+30.68%	+50.43%
FTB			
FTB	151.77	+40.83%	+62.98%
Non-FTB	97.70	+31.10%	-0.67%
Location			
Urban	114.20	+33.89%	+26.05%
Rural	104.62	+33.02%	-19.46%

Note: Second column of the table reports the estimated average value of information as defined in equation (3.13) in GBP per month. The third and fourth columns report the average percentage change in prices and search expenditures, respectively. Calculations made by simulating new prices and search behaviour from the new equilibrium, assuming that lenders drew had the same cost draws as in the baseline scenario.

changes in the total cost of search. Interestingly, in some cases the cost of search is lower in the baseline scenario. This is the case if in the new equilibrium most consumers receive only one offer, which is assumed to be costless. Typically this would be the case with getting a mortgage from one's home bank³⁷, though, unlike Allen et al. (2017), we do not have any additional data to corroborate this conjecture.

Disaggregating the results by product type (table 3.8) yields more interesting results. Remarkably, we find that not everyone benefits from intermediation. Borrowers with 3- and 5-year fixed rate deals those who amortise over less than 20 years in total either benefit by very little or are worse off. These results are driven by a modest changes in equilibrium prices which come hand in hand with massive reductions of total search expenditure. This implies that currently the level of consumer search for those products is low, and consequently, even with brokers present, commissions, market power and prices are already at high levels.³⁸

Moreover, the presence of intermediaries substantially affects pricing of mortgages with longer amortisation periods. Removing brokers from the markets would double the prices of 30-year and longer mortgages. Similarly, brokers help buyers of loans with flexible repayment schemes or cashback, both by exerting a negative pressure on lenders' prices and reducing overall search expenditure.

Overall, our results should be interpreted with some important caveats, as our model does not deliver predictions for some general equilibrium effects. First, we treated broker fees as exogenously given, and while based on our main result it would be tempting to conclude that increasing

³⁷The bank with which the borrower has her current/savings account.

³⁸Relatively higher mean markup estimates in 3.6 confirm this hypothesis.

Table 3.8: Value of information: breakdown by loan characteristics.

	VOI	% Δp	% ΔSE
Overall	112.15	+33.70%	+16.33%
LTV			
≤70	81.31	+27.02%	-26.37%
71-75	150.53	+48.41%	+90.52%
76-80	81.09	+26.08%	+26.06%
81-85	178.54	+53.77%	+98.04%
86-90	165.72	+38.28%	+48.36%
91-95	61.44	+12.08%	-50.29%
Deal			
2 years	165.13	+49.28%	+48.25%
3 years	-1.40	+1.47%	-67.62%
5 years	-3.99	-0.54%	-52.44%
Term			
≤10 years	4.89	+4.73%	-86.19%
(10;15]	-14.32	-1.99%	-77.46%
(15;20]	2.35	+2.52%	-69.37%
(20;25]	74.63	+23.76%	-43.84%
(25;30]	123.01	+36.02%	+7.88%
(30;35]	363.74	+104.36%	+284.55%
Value			
Q1	115.35	+34.79%	-50.97%
Q2	78.82	+23.13%	-28.31%
Q3	117.31	+34.19%	+30.91%
Q4	135.29	+42.05%	+112.17%
Flexible			
Flexible	17.26	+8.52%	-82.27%
Regular	126.93	+37.62%	+31.69%
Cashback			
No cashback	121.73	+36.72%	+29.33%
Cashback	48.35	+16.61%	-70.04%

Note: Second column of the table reports the estimated average value of information as defined in equation (3.13) in GBP per month. The third and fourth columns report the average percentage change in prices and search expenditures, respectively. Calculations made by simulating new prices and search behaviour from the new equilibrium, assuming that lenders drew had the same cost draws as in the baseline scenario.

them on average by over £100 per month would still make borrowers better off than in a hypothetical world without intermediation, doing that would drastically reduce the demand for broker services and force a lot of brokers to exit the market. Secondly, we did not allow for switching to a different mortgage type in the counterfactual. For example, knowing that brokers almost exclusively provide value to mortgagors shopping for 2-year fixed rate deals, it would be reasonable to assume that consumers would switch to 3- and 5-year fixed products. Finally, our model does not provide an estimate for total broker payoffs³⁹, so we do not attempt a full welfare analysis. With all that in mind, the result still can be interpreted in terms of value of information provided by brokers to the borrowers under the current market structure.

³⁹Woodward and Hall (2012) argue that brokers are indifferent between the main source of compensation (i.e. the contributions of the lender and the borrower) and only care about the sum of the two components. In our model, procurement fees, provided they are passed onto the borrowers, can be seen as a part of the estimated lenders' cost. In section C.1.4 we provide a robustness check where we adjust the estimated cost distributions by potential savings faced by the lenders assuming full pass-through of procurement fees. Overall, the average VOI drops from £112 to £97 and all the results are quantitatively similar to the ones presented here.

3.6.2 Market centralisation

In the second experiment we consider a hypothetical market centralisation. In the last two years, startups such as Habito⁴⁰ began to facilitate mortgage search by creating a free online platform propelled by machine learning algorithms which helps borrowers find the best offer given their needs and characteristics (so called *robo-advice*). Differently from traditional price comparison services, such as Moneyfacts, Habito does not just list prices but acts as an online broker who helps borrowers through the application process.

The counterfactual exercise in this section simulates the effects of extending such a technology to the entire market. This includes not allowing lenders to offer their products directly, only through the platform. In a centralised market, lenders price according to the standard first-price procurement bid formula:

$$\beta(c|\mathbf{x}^G, \mathbf{x}^H) = \beta(c|\mathbf{x}^H) = c + \frac{\int_{s=c}^{\bar{c}} (1 - \mathcal{H}(s|\mathbf{x}^H))^{J-1} ds}{(1 - \mathcal{H}(c|\mathbf{x}^H))^{J-1}} \quad [3.14]$$

Canonical results from auction theory Milgrom and Weber (1985) assure that the symmetric equilibrium of the bidding game is unique. Therefore, solving for the counterfactual is relatively straightforward as it only involves finding the best responses defined by 3.14 without having to determine optimal consumer search behaviour.

We look at projected benefits from such a market regulation assuming that the platform would be costless to access, completely frictionless environment. The results are summarised in table 3.9.

In a market without search, consumers would be paying on average £21.19 less per month (6.41% less than currently). The benefits would be further compounded by search expenditure savings of roughly £6.44. The sum of these two numbers corresponds to the average total increase in consumer surplus and the estimate is quantitatively very close to the result of a similar exercise conducted by Allen et al. (2017) who found that eliminating search frictions and limiting banks' market power in the Canadian market would increase consumer surplus by \$27.92.

As in the first counterfactual, the magnitude of the change varies across borrower and product types. Centralisation would be more beneficial for high income and older borrowers than for younger consumers with income below the median. In terms of product characteristics, 3- and 5-year fixed rate deals would become substantially cheaper (15 and 11%, respectively). Borrowers with higher LTV would on average benefit less from lower prices and lack of frictions.

⁴⁰For the of description of Habito's business model see e.g. The Financial Times: <https://www.ftadviser.com/mortgages/2017/01/24/habito-secures-5-5m-to-create-mortgage-platform/>

Table 3.9: Price changes in a centralised market.

	Δp	$\% \Delta p$	ΔSE		Δp	$\% \Delta p$	ΔSE
Overall	-21.19	-6.41%	-6.44	LTV			
Age				≤ 70	-31.48	-10.07%	-7.37
<30	-8.44	-1.99%	-6.11	71-75	-6.35	-1.73%	-4.57
30+	-23.95	-7.36%	-6.50	76-80	-20.17	-5.97%	-5.11
Income				81-85	-9.42	-2.34%	-5.30
Low	-8.19	-1.83%	-6.27	86-90	-12.26	-2.61%	-6.61
High	-25.94	-8.08%	-6.49	91-95	-5.99	-1.06%	-5.96
FTB				Deal			
FTB	-23.60	-6.74%	-8.06	2 years	-12.43	-4.12%	-7.29
Non-FTB	-18.85	-6.08%	-4.84	3 years	-50.26	-15.03%	-7.34
Location				5 years	-39.59	-11.11%	-4.33
Urban	-20.17	-6.07%	-6.70	Value			
Rural	-24.95	-7.67%	-5.47	Q1	-36.67	-10.69%	-11.07
Term				Q2	-21.14	-6.22%	-6.26
≤ 10 years	-24.95	-16.71%	-10.59	Q3	-14.31	-4.46%	-4.71
(10;15]	-53.62	-15.13%	-8.70	Q4	-12.18	-4.12%	-3.56
(15;20]	-47.88	-11.71%	-6.95	Flexible			
(20;25]	-37.52	-7.26%	-6.02	Flexible	-11.28	-3.75%	-7.91
(25;30]	-14.18	-3.98%	-5.47	Regular	-22.74	-6.82%	-6.21
(30;35]	-17.09	-5.25%	-5.02	Cashback			
				No cashback	-19.99	-6.01%	-6.36
				Cashback	-29.22	-9.08%	-6.92

Note: The second column of each panel shows the average absolute difference between prices charged by lenders in a centralised market and prices observed in the data. The third column is the same difference but in relative terms. The fourth and final column shows the average search expenditure savings a fully frictionless market would generate.

Establishing a market-wide platform would stimulate competition between lenders and make borrowers better off. However, in our framework, mortgagors would be the only side of the market that would benefit from such a regulation. In the following section we look at the effects centralisation would have for banks' markups. However, to comprehensively assess the cost of the regulation, we would need to take a stance on the profits of brokers and the potential sunk costs they would be facing if they had to exit the market.


3.6.3 Summary of findings

To summarise the findings from the two counterfactual exercises we conducted, we take another look at some aggregate statistics. Clearly, the current market structure lies between the two extreme cases considered in the counterfactuals. Since over 70% of mortgages are currently brokered, it is more similar to a centralised market. Figure 3.5 displays the distributions of markups and prices in the data and the two counterfactual scenarios.

Without brokers, lenders would enjoy substantially more market power. The average Lerner index would increase to 27.94% and the dispersion would also be much more substantial with 25% of borrowers being facing 40% or higher margins. In a centralised market, the median PCM is only 3.68% and the average price would decrease by 6.4%.

We conclude that the market is currently much more competitive than it would be if brokers were not present. While complete centralisation would reduce mortgage prices and lenders' mar-

Figure 3.5: Distributions of prices and price-cost margins in the counterfactual scenarios.



Note: Kernel estimates of the density of price-cost margins defined as $PCM_{ij} = \frac{p_{ij} - c_{ij}}{p_{ij}}$ (left panel) and estimated CDFs of prices (right panel).

gins, the overall change is would be modest and might not be sufficient to compensate for the (potentially high) costs of establishing such a platform. Therefore, we believe that the regulator should instead facilitate broker competition, allowing for new entry, but without necessarily banning lenders from direct sales.

3.7 Conclusion

Motivated by the observation that over 70% of mortgages in the UK are originated through intermediaries, this work attempts to estimate the value of information provided by brokers using a structural model of borrower search. Using the data on the universe of all originated mortgages in 2016 and 2017, we documented the existence of price dispersion and modest pecuniary benefits from using a broker. Furthermore, a large part of the decision whether to use a broker cannot be explained by observable borrower characteristics, which led us to a conclusion that the decision is driven by heterogeneity in unobserved costs of shopping around and filling out mortgage applications. Our main identifying assumption was that borrowers with high search cost, instead of contacting lenders directly, decide to use brokers select the best available deal on their behalf.

Our structural model nonparametrically identifies the distribution of search costs and lenders' costs of providing the loan. We estimate those primitives using techniques recently developed in the consumer search literature which leverage methods used for nonparametric auction estimation. We find that search costs can be substantial and are heterogeneous across different consumer types. On the supply side, the market appears to be relatively competitive, with average markups around

10%.

We use the estimates to simulate the effects of removing intermediaries from the market. The difference in consumer surplus is what we label as value of information. On average, we find that brokers advice is worth around £112 pounds per month, though not every borrower benefits from their presence. In the absence of brokers, firms would be enjoying significantly higher monopoly power and consumers would have to spend more on search. In the second counterfactual, we simulate the effects a hypothetical market centralisation, finding that it would lead only to a modest reduction of prices and lenders' market power.

This research makes two main contributions: first, the empirical results contribute to the policy discussion on the regulation of banks, brokers, and the mortgage market itself; secondly, the structural model presented here is relatively straightforward to estimate and simulate and the results are robust to distributional assumptions, and can be used to study a wider array of industries, where some consumers can access a platform while others purchase the good directly. Therefore we see it as an attractive framework for empirical studies of welfare effects of multisided platforms.

In the future, we wish to attempt to use recent results on the estimation of auctions with unobserved heterogeneity (Haile and Kitamura, 2018) to introduce broker heterogeneity into the model. This would allow us to relax the assumption that all intermediaries act as unbiased auctioneers and introduce potentially distorted advice. On the lender side, we could model the decision of offering the product via an intermediary using the results on estimating auctions with endogenous entry.

Concluding remarks

The papers comprising this thesis add to the empirical IO literature on oligopoly pricing and competition in markets with supply and demand frictions. The methodological contributions include the introduction of two new structural models of demand and price competition. In the future, I aim to extend the dynamic pricing model from chapter 1 to allow for persistent consumer heterogeneity and also adapt it to situations where the good is storable and consumers forward-looking. The search model introduced in chapter 2 can be further extended to allow for asymmetric firms. In an ongoing work we are also working on establishing methods of statistical testing between our model with heterogeneous firms and the standard model of Burdett and Judd (1983) with identical firms and complete information.

The main empirical contributions in chapter 1 included providing an estimate of hidden promotional costs in the vertical relationship between manufacturers and retailers and assessing their impact on consumer welfare. We also highlighted the interplay between consumer loyalty and adjustment costs which has not been properly studied in the literature so far. Future research can try to further investigate the issue of costly price adjustment in different industries, especially if the source of the cost is different than here, that is unrelated to promotional activity. The model proposed here, under a different set of identifying restrictions, can also be used in such settings. Chapter 3 suggested that the existence of intermediaries has a positive effect on competition and increases consumer surplus. However, the question whether middlemen are always welfare improving remains open in many other industries where they play a prominent role, including insurance, advertising and financial services. The structural approach proposed in this chapter can be used in the future to gain new insights about the role of intermediaries in those industries. Finally, the same approach can be used to analyse market where consumers have a choice to search directly amongst different sellers for the lowest price or use a platform aggregating different offers, for example when analysing demand for hotel services or airline tickets.

Appendix A

Supplementary material for chapter 1

A.1 Alternative demand specification

In this section of the appendix we present an alternative way of defining consumer loyalty in the demand model, in the vein of Dubé et al. (2008, 2009) and Pavlidis and Ellickson (2017). As opposed to the specification presented in section 1.5.3, aggregation here does not deliver a first-order Markov process on market shares. Instead, current period's demand realisation can be predicted using information on the fraction of consumers loyal to each of the products in the preceding period.

Rewriting (1.6):

$$u_{jt}^h = \delta_j - \eta \cdot p_{jt} + \gamma \cdot \mathbf{1}(q_t^h = j) + \xi_{jt}^h \quad [\text{A.1.1}]$$

To emphasise the difference between (1.6) and (A.1.1), we replace y_{t-1}^h with q_t^h . This variable indicates which good was purchased on a previous purchase occasion. In other words, if a consumer purchased good j in period 1 and chose the outside option in period 2, at the beginning of period 3 he will still be considered loyal to good j . Let $\mathbf{q}_t = [q_{1t}, \dots, q_{|\mathcal{J}|t}]'$ be an aggregate state variable, collecting the fractions of consumers loyal to each of the goods at the beginning of period t . Note that $\sum_{g=1}^{|\mathcal{J}|} q_{gt} = 1$, so the dimension of this state variable is $|\mathcal{J}| - 1$, which is 1 dimension less than \mathbf{s}_{t-1} . Aggregate market share of good j is now:

$$s_{jt}(a_{it}, \mathbf{a}_{-it}, \mathbf{q}_t) = \sum_{g=1}^{|\mathcal{J}|} q_{g,t} \cdot \Pr_t(j | \mathbf{p}_t(\mathbf{a}_t), q_t = g), \quad [\text{A.1.2}]$$

and the law of motion for the loyalty state is $\mathbf{q}_{t+1} = \Psi_t \mathbf{q}_t$, where $\Psi_t = \{\psi_{g \rightarrow j}\}_{g,j}^{|\mathcal{J}|}$ is a $|\mathcal{J}| \times |\mathcal{J}|$ transition matrix, which entry in row g , column j is:

$$\psi_{g \rightarrow j} = \begin{cases} \Pr_t(j|\mathbf{p}_t(\mathbf{a}_t), q_t = g) + \Pr_t(0|\mathbf{p}_t(\mathbf{a}_t), q_t = g) & \text{if } g = j \\ \Pr_t(j|\mathbf{p}_t(\mathbf{a}_t), q_t = g) & \text{if } g \neq j \end{cases} \quad [\text{A.1.3}]$$

This reformulation of the demand model implies a slight amendment to the firms' problem. Following the notation we introduced in section 1.5.2, the publicly observed vector of state variables is now: $\mathbf{z}_t = \{\mathbf{q}_t, \mathbf{a}_{t-1}\}$.

With multiple consumer types, one has to keep track of \mathbf{q}_t^h , that is the vector of loyalty states of type- h consumers. Assuming that the types are exogenous and fixed over time, one arrives at the model of Dubé et al. (2009). In general it is also possible to allow consumers to migrate between types, but these transitions have to be identified off the data without imposing further structure.

A.2 Main identification result

In this appendix we present our main identification result. To make it self-contained, we will repeat some of the notational assumptions we have been making throughout the main body of chapter 1. Also, to make the exposition clearer, we will be referring to a specific number of players, actions and cardinality of the set of possible market shares which will be the same as in our empirical application. This is without loss of generality and can be easily adapted to applications with different number of actions, players or alternative ways of discretising the state space.

Preliminaries

There are three players, producing two products each (four actions per player). There is also a generic good that can be chosen by consumers, but its price is exogenously given (hence there are 7 lagged market shares to keep track of). The vector of publicly observed state variables is $\mathbf{z}_t = (\mathbf{s}_{t-1}, \mathbf{a}_{t-1})$. We discretise last period's market shares into 3 bins, therefore the dimension of the state space \mathcal{Z} is: $|\mathcal{Z}| = 4^3 \cdot 3^7 = 64 \cdot 2187 = 139,968$. For simplicity we will refer to the action $(p_{i_1}^H, p_{i_2}^H)$ as HH . The payoff function of player i is:

$$\begin{aligned} \Pi(\mathbf{a}_t, \mathbf{z}_t, \varepsilon_{it}) &= \pi_i(a_{it}, \mathbf{a}_{-it}, \mathbf{s}_{t-1}) + \sum_{\ell \in \mathcal{A}_i} \varepsilon_{it}(\ell) \cdot \mathbf{1}(a_{it} = \ell) \\ &+ \sum_{\ell \in \mathcal{A}_i} \sum_{\ell' \neq \ell} SC_i^{\ell' \rightarrow \ell} \cdot \mathbf{1}(a_{it} = \ell, a_{i,t-1} = \ell'), \end{aligned} \quad [\text{A.2.1}]$$

Derivation

The non-stochastic dynamic payoff from choosing $a_{it} = \ell$ is:

$$\begin{aligned} \bar{v}_i(\ell, \mathbf{z}_t) &= \sum_{\substack{\mathbf{a}_{-it} \in \times_{j \neq i} \mathcal{A}_j}} \sigma_i(\mathbf{a}_{-it} | \mathbf{z}_t) \left[\pi_i(\ell, \mathbf{a}_{-it}, \mathbf{s}_{t-1}) + \beta \sum_{\mathbf{z}_{t+1}} G(\mathbf{z}_{t+1} | \mathbf{s}_{t-1}, \ell, \mathbf{a}_{-it}) \right. \\ &\quad \left. \cdot \underbrace{\int V_i(\mathbf{z}_{t+1}, \varepsilon_{t+1}) dQ(\varepsilon_{i,t+1})}_{\tilde{V}(\mathbf{z}_{t+1})} \right] + \sum_{\ell' \neq \ell} SC_i^{\ell' \rightarrow \ell} \cdot \mathbf{1}(a_{i,t-1} = \ell') \end{aligned} \quad [\text{A.2.2}]$$

Defining the differences with respect to the reference action HH we have:

$$\begin{aligned} \Delta \bar{v}_i(\ell, \mathbf{z}_t) &= \bar{v}_i(\ell, \mathbf{z}_t) - \bar{v}_i(HH, \mathbf{z}_t) \\ &= \sum_{\substack{\mathbf{a}_{-it} \in \times_{j \neq i} \mathcal{A}_j}} \sigma_i(\mathbf{a}_{-it} | \mathbf{z}_t) \left\{ \underbrace{\pi_i(\ell, \mathbf{a}_{-it}, \mathbf{s}_{t-1}) - \pi_i(HH, \mathbf{a}_{-it}, \mathbf{s}_{t-1})}_{\Delta \pi_i^\ell(\mathbf{a}_{-it}, \mathbf{s}_{t-1})} \right\} \\ &\quad + \sum_{\substack{\mathbf{a}_{-it} \in \times_{j \neq i} \mathcal{A}_j}} \sigma_i(\mathbf{a}_{-it} | \mathbf{z}_t) \left\{ \beta \sum_{\mathbf{z}_{t+1}} \underbrace{[G(\mathbf{z}_{t+1} | \mathbf{s}_{t-1}, \ell, \mathbf{a}_{-it}) - G(\mathbf{z}_{t+1} | \mathbf{s}_{t-1}, HH, \mathbf{a}_{-it})]}_{\Delta G^\ell(\mathbf{z}_{t+1} | \mathbf{a}_{-it}, \mathbf{s}_{t-1})} \right\} \tilde{V}(\mathbf{z}_{t+1}) \\ &\quad + \underbrace{\sum_{\ell' \neq \ell} [SC_i^{\ell' \rightarrow \ell} \cdot \mathbf{1}(a_{i,t-1} = \ell') - SC_i^{\ell' \rightarrow HH} \cdot \mathbf{1}(a_{i,t-1} = \ell')]}_{\Delta SC_i^\ell(a_{i,t-1})} \end{aligned}$$

Using the newly introduced notation, we have:

$$\begin{aligned} \Delta \bar{v}_i(\ell, \mathbf{z}_t) &= \sum_{\substack{\mathbf{a}_{-it} \in \times_{j \neq i} \mathcal{A}_j}} \sigma_i(\mathbf{a}_{-it} | \mathbf{z}_t) \left\{ \underbrace{\Delta \pi_i^\ell(\mathbf{a}_{-it}, \mathbf{s}_{t-1}) + \beta \sum_{\mathbf{z}_{t+1}} \Delta G^\ell(\mathbf{z}_{t+1} | \mathbf{a}_{-it}, \mathbf{s}_{t-1}) \tilde{V}(\mathbf{z}_{t+1})}_{\lambda_i(\ell, \mathbf{a}_{-it}, \mathbf{s}_{t-1})} \right\} \\ &\quad + \Delta SC_i^\ell(a_{i,t-1}) \end{aligned} \quad [\text{A.2.3}]$$

Thinking back about the dimension of the problem, for each of the three remaining (that is, excluding HH) actions of player i , there are $4^2 \cdot 3^7 = 16 \cdot 2187 = 34992$ $\lambda_i(\ell, *)$ terms. Rewriting (A.2.3) in vector form:

$$\Delta \bar{v}_i(\ell, \mathbf{z}_t) = \boldsymbol{\sigma}_i(\mathbf{z}_t)' \boldsymbol{\lambda}_i(\ell, \mathbf{s}_{t-1}) + \Delta SC_i^\ell(a_{i,t-1}), \quad [\text{A.2.4}]$$

where $\boldsymbol{\sigma}_i(\mathbf{z}_t) = [\sigma_i(\mathbf{a}_{-it} | \mathbf{z}_t)]_{\mathbf{a}_{-it}}$ and $\boldsymbol{\lambda}_i(\ell, \mathbf{s}_{t-1}) = [\lambda_i(\ell, \mathbf{a}_{-it}, \mathbf{s}_{t-1})]_{\mathbf{a}_{-it}}$ are 16×1 column vectors. (A.2.4) holds for all of the 139,968 points in the state space. To make things more explicit, use the fact that \mathbf{z}_t can be partitioned into $(\mathbf{a}_{t-1}, \mathbf{s}_{t-1})$. Furthermore:

$$\begin{aligned} \mathbf{a}_{t-1} &= \{\mathbf{a}_{t-1}^1, \mathbf{a}_{t-1}^2, \dots, \mathbf{a}_{t-1}^{64}\} \\ \mathbf{s}_{t-1} &= \{\mathbf{s}_{t-1}^1, \mathbf{s}_{t-1}^2, \dots, \mathbf{s}_{t-1}^{2187}\} \end{aligned}$$

For \mathbf{s}_{t-1}^1 the system can be written as:

$$\begin{cases} \Delta \bar{v}_i(\ell, \mathbf{a}_{t-1}^1, \mathbf{s}_{t-1}^1) = \boldsymbol{\sigma}_i(\mathbf{a}_{t-1}^1, \mathbf{s}_{t-1}^1)' \boldsymbol{\lambda}_i(\ell, \mathbf{s}_{t-1}^1) + \Delta SC_i^\ell(\mathbf{a}_{t-1}^1) \\ \Delta \bar{v}_i(\ell, \mathbf{a}_{t-1}^2, \mathbf{s}_{t-1}^1) = \boldsymbol{\sigma}_i(\mathbf{a}_{t-1}^2, \mathbf{s}_{t-1}^1)' \boldsymbol{\lambda}_i(\ell, \mathbf{s}_{t-1}^1) + \Delta SC_i^\ell(\mathbf{a}_{t-1}^2) \\ \vdots \\ \Delta \bar{v}_i(\ell, \mathbf{a}_{t-1}^{64}, \mathbf{s}_{t-1}^1) = \boldsymbol{\sigma}_i(\mathbf{a}_{t-1}^{64}, \mathbf{s}_{t-1}^1)' \boldsymbol{\lambda}_i(\ell, \mathbf{s}_{t-1}^1) + \Delta SC_i^\ell(\mathbf{a}_{t-1}^{64}) \end{cases}$$

Vectorizing again:

$$\Delta \bar{\mathbf{v}}_i(\ell, \mathbf{s}_{t-1}^1) = \boldsymbol{\sigma}_i(\mathbf{s}_{t-1}^1) \boldsymbol{\lambda}_i(\ell, \mathbf{s}_{t-1}^1) + \Delta \mathbf{S} \mathbf{C}_i^\ell, \quad [\text{A.2.5}]$$

where $\bar{\mathbf{v}}_i(\ell, \mathbf{s}_{t-1}^1) = [\Delta \bar{v}_i(\ell, \mathbf{a}_{t-1}, \mathbf{s}_{t-1}^1)]_{\mathbf{a}_{t-1}}$ is a 64×1 vector, $\boldsymbol{\sigma}_i(\mathbf{s}_{t-1}^1) = [\boldsymbol{\sigma}_i(\mathbf{a}_{t-1}, \mathbf{s}_{t-1}^1)']_{\mathbf{a}_{t-1}}$ is a 64×16 matrix and $\Delta \mathbf{S} \mathbf{C}_i^\ell = [\Delta SC_i^\ell(\mathbf{a}_{t-1})]_{\mathbf{a}_{t-1}}$ is a 64×1 vector. In matrix notation, for all \mathbf{s}_{t-1} , this becomes:

$$\Delta \bar{\mathbf{v}}_i(\ell) = \underbrace{\begin{bmatrix} \boldsymbol{\sigma}_i(\mathbf{s}_{t-1}^1) & \mathbf{0} \\ & \ddots \\ \mathbf{0} & \boldsymbol{\sigma}_i(\mathbf{s}_{t-1}^{2187}) \end{bmatrix}}_{(2187 \cdot 64) \times (2187 \cdot 16)} \underbrace{\begin{bmatrix} \boldsymbol{\lambda}_i(\ell, \mathbf{s}_{t-1}^1) \\ \vdots \\ \boldsymbol{\lambda}_i(\ell, \mathbf{s}_{t-1}^{2187}) \end{bmatrix}}_{(2187 \cdot 16) \times 1} + \widetilde{\Delta \mathbf{S} \mathbf{C}_i^\ell} \quad [\text{A.2.6}]$$

We will be referring to the block-diagonal matrix containing player i 's beliefs as $\boldsymbol{\sigma}$. It can be written more compactly as a Kronecker product of an identity matrix I and matrix containing beliefs:

$$\begin{aligned} \Delta \bar{\mathbf{v}}_i(\ell) &= \begin{bmatrix} I_{2187} \otimes \begin{bmatrix} \boldsymbol{\sigma}_i(\mathbf{s}_{t-1}^1) \\ \vdots \\ \boldsymbol{\sigma}_i(\mathbf{s}_{t-1}^{2187}) \end{bmatrix} \end{bmatrix} \begin{bmatrix} \boldsymbol{\lambda}_i(\ell, \mathbf{s}_{t-1}^1) \\ \vdots \\ \boldsymbol{\lambda}_i(\ell, \mathbf{s}_{t-1}^{2187}) \end{bmatrix} + \widetilde{\Delta \mathbf{S} \mathbf{C}_i^\ell} \\ &= \boldsymbol{\sigma}_i \boldsymbol{\lambda}_i(\ell) + \widetilde{\Delta \mathbf{S} \mathbf{C}_i^\ell} \end{aligned}$$

Everything we showed so far was for a selected action $\ell \in \mathcal{A}_i \setminus \{HH\}$. We can now define $\Delta \bar{\mathbf{v}}_i = [\bar{\mathbf{v}}_i(HL); \bar{\mathbf{v}}_i(LH); \bar{\mathbf{v}}_i(LL)]'$, so that:

$$\begin{aligned} \Delta \bar{\mathbf{v}}_i &= [I_3 \otimes \boldsymbol{\sigma}_i] \begin{bmatrix} \boldsymbol{\lambda}_i(HL) \\ \boldsymbol{\lambda}_i(LH) \\ \boldsymbol{\lambda}_i(LL) \end{bmatrix} + \begin{bmatrix} \widetilde{\Delta \mathbf{S} \mathbf{C}_i^{HL}} \\ \widetilde{\Delta \mathbf{S} \mathbf{C}_i^{LH}} \\ \widetilde{\Delta \mathbf{S} \mathbf{C}_i^{LL}} \end{bmatrix} \\ &= \mathbf{Z}_i \boldsymbol{\lambda}_i + \widetilde{\Delta \mathbf{S} \mathbf{C}_i} \end{aligned} \quad [\text{A.2.7}]$$

The dimension of the object on the LHS of (A.2.7) is $(139968 \cdot 3 \times 1) = 419904 \times 1$. Define the following 419904×419904 projection matrix:

$$\mathbf{M}_i^Z = I_{419904} - \mathbf{Z}_i(\mathbf{Z}_i' \mathbf{Z}_i)^{-1} \mathbf{Z}_i' \quad [\text{A.2.8}]$$

So far we have not discussed $\widetilde{\Delta \mathbf{S} \mathbf{C}}_i$ in detail, but it can be written as: $\widetilde{\Delta \mathbf{S} \mathbf{C}}_i = \widetilde{\mathbf{D}}_i \Delta \mathbf{S} \mathbf{C}_i$ where $\widetilde{\mathbf{D}}_i$ is a $419904 \times \kappa_i$ matrix of zeros and ones which are a natural consequence of the indicator functions used while defining the profit function. κ_i is the number of dynamic parameters to estimate for player i and $\Delta \mathbf{S} \mathbf{C}_i$ is a $\kappa_i \times 1$ vector of parameters to identify. Multiplying both sides of (A.2.7) by the projection matrix defined in (A.2.8), we have:

$$\begin{aligned} \mathbf{M}_i^Z \Delta \bar{\mathbf{v}}_i &= \mathbf{M}_i^Z \widetilde{\mathbf{D}}_i \Delta \mathbf{S} \mathbf{C}_i \\ \widetilde{\mathbf{D}}_i' \mathbf{M}_i^Z \Delta \bar{\mathbf{v}}_i &= \widetilde{\mathbf{D}}_i' \mathbf{M}_i^Z \widetilde{\mathbf{D}}_i \Delta \mathbf{S} \mathbf{C}_i \\ \Delta \mathbf{S} \mathbf{C}_i &= (\widetilde{\mathbf{D}}_i' \mathbf{M}_i^Z \widetilde{\mathbf{D}}_i)^{-1} (\widetilde{\mathbf{D}}_i' \mathbf{M}_i^Z \Delta \bar{\mathbf{v}}_i) \end{aligned} \quad [\text{A.2.9}]$$

(A.2.9) defines the identifying correspondence for player i . We can proceed in an identical fashion to recover the parameters for the remaining players. There is also a straightforward way to incorporate equality restrictions across players an estimate $\{\Delta \mathbf{S} \mathbf{C}_i\}_{i=1}^N$ for all players in one step.

Computation

The main computational challenge here lies in the construction of the projection matrix \mathbf{M}_i^Z which involves inverting the matrix $\mathbf{Z}_i' \mathbf{Z}_i$ of size $3 \cdot 34992 \times 3 \cdot 34992$. However, a closer inspection reveals that this matrix is block-diagonal. To see this, rewrite \mathbf{Z}_i :

$$\mathbf{Z}_i = \begin{bmatrix} \begin{bmatrix} \sigma_i(\mathbf{s}_{t-1}^1) & & \mathbf{0} \\ & \ddots & \\ \mathbf{0} & & \sigma_i(\mathbf{s}_{t-1}^{2187}) \end{bmatrix} & & & \mathbf{0} \\ & & \begin{bmatrix} \sigma_i(\mathbf{s}_{t-1}^1) & & \mathbf{0} \\ & \ddots & \\ \mathbf{0} & & \sigma_i(\mathbf{s}_{t-1}^{2187}) \end{bmatrix} & & \\ & & & & \begin{bmatrix} \sigma_i(\mathbf{s}_{t-1}^1) & & \mathbf{0} \\ & \ddots & \\ \mathbf{0} & & \sigma_i(\mathbf{s}_{t-1}^{2187}) \end{bmatrix} \\ & \mathbf{0} & & & \end{bmatrix}$$

Recall that each of the $\sigma_i(\cdot)$'s is a 64×16 matrix. Multiplying \mathbf{Z}_i by its transpose we have:

$$\mathbf{Z}'_i \mathbf{Z}_i = \begin{bmatrix} \left[\begin{array}{cc} \sigma_i(\mathbf{s}_{t-1}^1)' \sigma_i(\mathbf{s}_{t-1}^1) & \mathbf{0} \\ \vdots & \vdots \\ \mathbf{0} & \sigma_i(\mathbf{s}_{t-1}^{2187})' \sigma_i(\mathbf{s}_{t-1}^{2187}) \end{array} \right] & \mathbf{0} \\ \mathbf{0} & \left[\begin{array}{c} [] \\ [] \end{array} \right] \end{bmatrix}$$

Now each of the $\sigma_i(\cdot)' \sigma_i(\cdot)$ entries is a 16×16 matrix, so in the end to obtain the inverse of $\mathbf{Z}'_i \mathbf{Z}_i$ we have to invert 2187 16×16 matrices, which in principle should be much faster and accurate than inverting one big matrix. In practice we can proceed as follows:

1. Construct 2187 projection matrices:

$$\mathbf{M}_i^{\mathbf{Z}}(\cdot) = I_{64} - \sigma_i(\cdot) [\sigma_i(\cdot)' \sigma_i(\cdot)]^{-1} \sigma_i(\cdot)'$$

2. Build the matrix \mathbf{M}_i^{λ}
3. Recover $\Delta \mathbf{SC}_i$

A.3 Identification of promotional costs

This appendix shows how assuming that adjustment costs are only paid by firms if they change prices from high to low allows to point-identify the vector of costs consisting of a separate parameters for each product. We start with assumptions R1-3:

Assumption (R1). *Adjustment costs are incurred only when switching from high to low price.*

Assumption (R2). *Adjustment cost associated with one product is independent of the current and lagged promotional status of other products.*

R2 is a natural assumption, and allows us to impose equality restriction across $\mathbf{a}_{-i,t-1}$ in the switching cost part of (1.2). Finally, consider the situation in which prices of more than one product of a firm move in the same direction. R3 says that we can express the cost of taking this action as a sum of individual price adjustments of the products involved:

Assumption (R3). *There are no economies of scope associated with price promotions on multiple products of the same firm.*

R1-2 will be sufficient to identify one cost of adjusting prices per product, and R3 can be just used to reduce the dimension of the parameter vector. The identifying power of our assumptions is summarised by the following proposition:

Proposition 1. *Under assumptions R1-2, the matrix $\tilde{\mathbf{D}}_i$ satisfies the requirements of theorem 2 in Komarova et al. (forthcoming) and for each player one can identify $|\mathcal{A}_i| - 1$ parameters in \mathbf{SC}_i . Adding assumption R3 reduces the number of parameters to $|\mathcal{J}_i|$.*

For clarity of exposition we prove proposition 1 for a two-product duopoly case. Generalising it to more players and products is straightforward.

Setup

Consider a simplified version of the model described in section 3.3: suppose there are **two** players which we denote as $i = \{a, b\}$ producing **two** differentiated goods each, whose sets we denote as $\mathcal{J}_i = \{i_1, i_2\}$.

Conditional on $(\mathbf{s}_{t-1}, \mathbf{a}_{t-1}, \varepsilon_{it})$, player i chooses an action a_{it} from the set \mathcal{A}_i to maximise the discounted sum of profits given her beliefs about the actions of the competitor. The decision variable in this game is the vector of prices of all goods manufactured by a player. Since prices are constrained to take only two values, regular (H) and sale (L), the cardinalities of both \mathcal{A}_a and \mathcal{A}_b are $2^{|\mathcal{J}_a|} = 2^{|\mathcal{J}_b|} = 4$.

Specifically $\mathcal{A}_a = \left\{ \underbrace{(p_{a_1}^H, p_{a_2}^H)}_{HH}; \underbrace{(p_{a_1}^H, p_{a_2}^L)}_{HL}; \underbrace{(p_{a_1}^L, p_{a_2}^H)}_{LH}; \underbrace{(p_{a_1}^L, p_{a_2}^L)}_{LL} \right\}$, where H/L denotes regular/sale price, \mathcal{A}_b is defined analogously.

This implies that without further restrictions there are 12 parameters per player:

$$\mathbf{SC}_i = [SC_i^{HH \rightarrow HH}, SC_i^{LH \rightarrow HH}, SC_i^{LL \rightarrow HH}, SC_i^{HH \rightarrow HL}, SC_i^{LH \rightarrow HL}, SC_i^{LL \rightarrow HL}, SC_i^{HH \rightarrow LH}, SC_i^{HL \rightarrow LH}, SC_i^{LL \rightarrow LH}, SC_i^{HH \rightarrow LL}, SC_i^{HL \rightarrow LL}, SC_i^{LH \rightarrow LL}]'$$

Under R1-2 there are three dynamic parameters to identify for each player, that is $SC_i^{HH \rightarrow LL}$, $SC_i^{HH \rightarrow HL}$, $SC_i^{HH \rightarrow LH}$, whereas R3 reduces this number to just two. With an arbitrary number of actions, $|\mathcal{A}_i|$ initially there are $|\mathcal{A}_i| \cdot (|\mathcal{A}_i| - 1)$ possible adjustment costs, $(|\mathcal{A}_i| - 1)$ under R1-2 and $|\mathcal{J}_i|$ under R1-3.

Identification

As previously, we take HH to be the reference action, so that: $\Delta \bar{v}_i(\ell, \mathbf{a}_{-it}, \mathbf{s}_{t-1}) \equiv \bar{v}_i(\ell, \mathbf{a}_{-it}, \mathbf{s}_{t-1}) - \bar{v}_i(HH, \mathbf{a}_{-it}, \mathbf{s}_{t-1})$. The reason why we use HH is that thanks to R1, no cost is ever incurred in

period t if $a_{it} = HH$. Therefore, for player a we have:

$$\Delta \bar{v}_a(\ell, \mathbf{a}_{-it}, \mathbf{s}_{t-1}) = \sum_{\mathbf{a}_{bt} \in \mathcal{A}_b} \Pr_b(\mathbf{a}_{bt} | \mathbf{a}_{t-1}, \mathbf{s}_{t-1}) \lambda_a(\ell, \mathbf{a}_{bt}, \mathbf{s}_{t-1}) + SC_a^{\ell' \rightarrow \ell} \cdot \mathbf{1}(a_{i,t-1} = \ell') \quad [\text{A.3.1}]$$

where $\lambda_a(\cdot)$ is defined as in (A.2.3). What remains to be verified is that the matrix of zeros and ones resulting from stacking the expressions (A.3.1) for all actions of player a and all possible values of the state variables is indeed of full rank and does not contain a column of ones.

To show this, we invoke lemma 5 in Komarova et al. (forthcoming) and write the expression for one possible realisation of lagged market shares, \mathbf{s}_{t-1} :

$$\Delta \bar{\mathbf{v}}_a(\mathbf{s}_{t-1}) = (I_{|\mathcal{A}_a|-1} \otimes \mathbf{Z}_a(\mathbf{s}_{t-1})) \boldsymbol{\lambda}_a(\mathbf{s}_{t-1}) + \tilde{\mathbf{D}}_a(\mathbf{s}_{t-1}) \mathbf{S} \mathbf{C}_a \quad [\text{A.3.2}]$$

where:

- $\Delta \bar{\mathbf{v}}_a(\mathbf{s}_{t-1}) = \{\Delta \bar{v}_a(a_{at}, \mathbf{a}_{t-1}, \mathbf{s}_{t-1})\}_{a_{at} \in \mathcal{A}_a \setminus \{HH\}}$ is a $(|\mathcal{A}_a| - 1) \cdot |\mathcal{A}_a| \cdot |\mathcal{A}_b| \times 1$ vector,
- $\mathbf{Z}_a(\mathbf{s}_{t-1}) = \{\Pr_b(a_{bt} | \mathbf{a}_{t-1}, \mathbf{s}_{t-1})\}_{(a_{bt}, \mathbf{a}_{t-1}) \in \mathcal{A}_b \times (\mathcal{A}_a \times \mathcal{A}_b)}$ is a $|\mathcal{A}_a| \cdot |\mathcal{A}_b| \times |\mathcal{A}_b|$ matrix,
- $\boldsymbol{\lambda}_a(\mathbf{s}_{t-1}) = \{\lambda_a(a_{at}, a_{bt}, \mathbf{s}_{t-1})\}_{(a_{at}, a_{bt}) \in (\mathcal{A}_a \setminus \{HH\}, \mathcal{A}_b)}$ is a $(|\mathcal{A}_a| - 1) \cdot |\mathcal{A}_b| \times 1$ vector,
- $\tilde{\mathbf{D}}_a(\mathbf{s}_{t-1})$ is a $(|\mathcal{A}_a| - 1) \cdot |\mathcal{A}_a| \cdot |\mathcal{A}_b| \times \kappa$ matrix,
- $\mathbf{S} \mathbf{C}_a$ is a $\kappa \times 1$ vector of parameters to identify.

To show the content of the objects in (A.3.2) we rewrite it as (A.3.4). For the sake of brevity, \mathbf{s}_{t-1} was dropped from the notation. We can immediately see from there that $\tilde{\mathbf{D}}_a(\mathbf{s}_{t-1})$ satisfies the rank condition. Imposing R3 only changes the last component of the sum on the RHS of (A.3.4), so that it becomes:

$$\tilde{\mathbf{D}}_a(\mathbf{s}_{t-1})\mathbf{S}\mathbf{C}_a = \begin{bmatrix} 1 & 0 \\ 1 & 0 \\ 1 & 0 \\ 1 & 0 \\ 1 & 0 \\ 0 & 0 \\ 0 & 0 \\ 0 & 0 \\ 0 & 0 \\ 0 & 0 \\ 1 & 0 \\ 1 & 0 \\ 1 & 0 \\ 1 & 0 \\ 0 & 0 \\ 0 & 0 \\ 0 & 0 \\ 0 & 0 \\ 0 & 0 \\ 0 & 0 \\ \hline 0 & 1 \\ 0 & 1 \\ 0 & 1 \\ 0 & 1 \\ 0 & 1 \\ 0 & 1 \\ 0 & 1 \\ 0 & 1 \\ 0 & 1 \\ 0 & 1 \\ 0 & 0 \\ 0 & 0 \\ 0 & 0 \\ 0 & 0 \\ 0 & 0 \\ 0 & 0 \\ 0 & 0 \\ 0 & 0 \\ 0 & 0 \\ 0 & 0 \\ \hline 1 & 1 \\ 1 & 1 \\ 1 & 1 \\ 1 & 1 \\ 0 & 1 \\ 0 & 1 \\ 0 & 1 \\ 0 & 1 \\ 1 & 0 \\ 1 & 0 \\ 1 & 0 \\ 1 & 0 \\ 0 & 0 \\ 0 & 0 \\ 0 & 0 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} S C_a^{HH \rightarrow HL} \\ S C_a^{HH \rightarrow LH} \end{bmatrix}$$

To arrive at (A.2.7) we vertically stack the vectors and matrices in (A.3.2) for all possible combinations of \mathbf{s}_{t-1} . But since $\tilde{\mathbf{D}}_a(\mathbf{s}_{t-1})$ does not vary across \mathbf{s}_{t-1} , then $\tilde{\mathbf{D}}_a = \mathbf{j} \otimes \tilde{\mathbf{D}}_a(\mathbf{s}_{t-1})$, where \mathbf{j} is a column vector of ones whose dimension is equal to the cardinality of \mathbf{s}_{t-1} .

We can now directly use the identifying correspondence below to recover the costs:

$$\begin{bmatrix} SC_a^{HH \rightarrow HL} \\ SC_a^{HH \rightarrow LH} \\ SC_a^{HH \rightarrow LL} \end{bmatrix} = (\tilde{\mathbf{D}}_a'(\mathbf{s}_{t-1}) \mathbf{M}_{I_{|\mathcal{A}_a|-1} \otimes \mathbf{Z}_a(\mathbf{s}_{t-1})} \tilde{\mathbf{D}}_a(\mathbf{s}_{t-1}))^{-1} \tilde{\mathbf{D}}_a'(\mathbf{s}_{t-1}) \mathbf{M}_{I_{|\mathcal{A}_a|-1} \otimes \mathbf{Z}_a(\mathbf{s}_{t-1})} \Delta \bar{\mathbf{v}}_a(\mathbf{s}_{t-1}),$$

[A.3.3]

discount factor.

Using the fact the state transitions in our model are deterministic – see equation (1.10) – we can write the *ex ante* value function in problem (1.4) as:

$$V_i(\mathbf{a}_{t-1}, \mathbf{s}_{t-1}) = \sum_{\mathbf{a}_t \in \mathcal{A}_i} \sigma_i(\mathbf{a}_t | \mathbf{a}_{t-1}, \mathbf{s}_{t-1}) \left\{ \tilde{\Pi}_i(\mathbf{a}_t, \mathbf{a}_{t-1}, \mathbf{s}_{t-1}) + \beta V_i(\mathbf{a}_t, \mathbf{s}(\mathbf{a}_t, \mathbf{s}_{t-1})) \right\}, \quad [\text{A.4.1}]$$

where $V_i(\mathbf{z}_{t+1}) = \int V_i(\mathbf{z}_{t+1}, \varepsilon_{t+1}) dQ(\varepsilon_{i,t+1})$ and $\tilde{\Pi}_i(\mathbf{a}_t, \mathbf{a}_{t-1}, \mathbf{s}_{t-1})$ is the (conditional) expectation of the payoff function $\Pi_i(\mathbf{a}_t, \mathbf{a}_{t-1}, \mathbf{s}_{t-1}, \varepsilon_{it}(a_{it}))$ with respect to ε_{it} when states are $(\mathbf{a}_{t-1}, \mathbf{s}_{t-1})$ and current actions are \mathbf{a}_t , and $\mathbf{s}(\mathbf{a}_t, \mathbf{s}_{t-1})$ is the vector of current shares – implied by equation (1.10) – when past shares are \mathbf{s}_{t-1} and current actions are \mathbf{a}_t . As in Sweeting (2013) we approximate $V_i(\mathbf{z}_t)$ using the following parametric function:

$$V_i(\mathbf{z}_t) \simeq \sum_{k=1}^K \lambda_{ki} \phi_{ki}(\mathbf{z}_t) \equiv \Phi_i(\mathbf{z}_t) \lambda_i, \quad [\text{A.4.2}]$$

where λ_{ki} is a coefficient and $\phi_{ki}(\cdot)$ is a well-defined function mapping the state vector into the set of real numbers. In our case, $\phi_{ki}(\cdot)$ are flexible functions of shares and prices of the firms. In practice, the variables we use to approximate the value functions include (i) (past) actions of all firms, (ii) second order polynomials of (past) shares of all products, (iii) interactions between (past) actions and shares of the different products and (iv) second order polynomials of the interactions between (past) actions and shares. We experimented with third and fourth order polynomials of shares and interactions between shares and actions but the results did not change significantly.

Notice that under this formulation solving for the value function requires that one computes only K parameters (λ_{ki} 's) for each manufacturer. By substituting this equation into the *ex ante* value function we can solve for $\lambda_i = [\lambda_{1i} \lambda_{2i} \dots \lambda_{Ki}]'$ in closed-form as a function of the primitives of the model, states and beliefs. Substituting (A.4.2) into (A.4.1) we get:

$$\Phi_i(\mathbf{a}_{t-1}, \mathbf{s}_{t-1}) \lambda_i = \sum_{\mathbf{a}_t \in \mathcal{A}} \sigma_i(\mathbf{a}_t | \mathbf{a}_{t-1}, \mathbf{s}_{t-1}) \left\{ \tilde{\Pi}_i(\mathbf{a}_t, \mathbf{a}_{t-1}, \mathbf{s}_{t-1}) + \beta \Phi_i(\mathbf{a}_t, \mathbf{s}(\mathbf{a}_t, \mathbf{s}_{t-1})) \lambda_i \right\}.$$

To simplify the notation let $\tilde{\Pi}_i^*(\mathbf{a}_{t-1}, \mathbf{s}_{t-1})$ and $\Phi_i^*(\mathbf{s}_{t-1})$ be the conditional expectations of $\tilde{\Pi}_i(\mathbf{a}_t, \mathbf{a}_{t-1}, \mathbf{s}_{t-1})$ and of $\Phi_i(\mathbf{a}_t, \mathbf{s}(\mathbf{a}_t, \mathbf{s}_{t-1}))$ with respect to current actions, respectively. Therefore, we can rewrite equation above as:

$$(\Phi_i(\mathbf{a}_{t-1}, \mathbf{s}_{t-1}) - \beta \Phi_i^*(\mathbf{s}_{t-1})) \lambda_i = \tilde{\Pi}_i^*(\mathbf{a}_{t-1}, \mathbf{s}_{t-1}).$$

Stacking this equation for every possible state in S we have that:

$$(\Phi_i - \beta\Phi_i^*) \lambda_i = \tilde{\Pi}_i^*,$$

where Φ_i and Φ_i^* are $N_s \times K$ matrices that depend on states and beliefs and $\tilde{\Pi}_i^*$ is a $N_s \times 1$ vector of expected profits that depends on state, beliefs and parameters, N_s being the number of states observed in the data. Assuming $K < N_s$, this expression can be rewritten as:

$$\lambda_i = \left[(\Phi_i - \beta\Phi_i^*)' (\Phi_i - \beta\Phi_i^*) \right]^{-1} \left[(\Phi_i - \beta\Phi_i^*)' \tilde{\Pi}_i^* \right]. \quad [\text{A.4.3}]$$

Inserting (A.4.3) into (A.4.2) we obtain the unconditional value functions associated to problem (1.4); given the logit assumption on ε_{it} we can calculate the probability of each action solving problem (1.4). Having estimated adjustment costs outside of the dynamic model and having calibrated H and marginal costs, the only parameter to be estimated inside the dynamic model is the discount factor. We do this by choosing the discount factor that minimises the difference between estimated action probabilities and the probabilities implied by the structural model, which are defined based on the approximation explained above (see Komarova et al. (forthcoming)).

A.5 Model solution

To solve the model we use an algorithm similar to that described in Sweeting (2013). The algorithm works as follows:

1. In step s we calculate $\lambda(\sigma^s)$ as a function of the vector of beliefs, σ^s , substituting equation (A.4.2) into the *ex ante* value function and solving for $\lambda = [\lambda_1 \lambda_2 \dots \lambda_k]$ in closed-form as a function of the primitives of the model, states and beliefs;
2. We use $\lambda(\sigma^s)$ to calculate choice specific value functions for each of the selected states and the multinomial logit formula implied by the model to update the vector of beliefs, $\tilde{\sigma}$;
3. If the value of the Euclidian norm $\|\sigma^s - \tilde{\sigma}\|$ is sufficiently small we stop the procedure and save $\tilde{\sigma}$ as the equilibrium vector of probabilities implied by the model, $\tilde{\sigma} = \sigma^*$; if $\|\sigma^s - \tilde{\sigma}\|$ is larger than the tolerance we update $\sigma^{s+1} = \psi\tilde{\sigma} + (1 - \psi)\sigma^s$, where ψ is a number between 0 and 1, and restart the procedure.

The tolerance used on $\|\sigma^s - \tilde{\sigma}\|$ was 10^{-3} and the value of ψ used to update σ^s to σ^{s+1} was 0.5. We have made several attempts using lower values for the tolerance on $\|\sigma^s - \tilde{\sigma}\|$ and for

ψ . All these attempts generated very similar equilibrium probabilities, but the time to achieve convergence was larger. The initial guess used to start the algorithm, σ^0 , is equal to the estimated CCPs evaluated at the corresponding state. To check the robustness of our results to changes in the initial guess we arbitrarily changed the original initial guess multiplying it by several factors between 0 and 1. For all our attempts the resulting equilibrium vector of probabilities was the same.

For the counterfactuals we have to simulate the model for states that are not observed in the data – i.e. we need estimates of σ^* for states that are not in the data. To do this we assumed that the solution of the model, σ^* , for the relevant counterfactual scenario is a logistic function of a linear index of states – i.e. the same function that we used to compute the CCPs. Mathematically, let $\sigma_i^*(a_i = k|\mathbf{z})$ be the probability that firm i plays $a_i = k$ when the state vector is \mathbf{z} . We assume that:

$$\sigma_i^*(a_i = k|\mathbf{z}) = \frac{\exp(\mathbf{z}'\gamma_k)}{\sum_{k'} \exp(\mathbf{z}'\gamma_{k'})}. \quad [\text{A.5.1}]$$

Dividing it by the probability of an anchor choice, say $a_i = HH$, normalising $\gamma_1 = 0$ and taking logs we have $\ln \{\sigma_i^*(a_i = k|\mathbf{z})\} - \ln \{\sigma_i^*(a_i = 1|\mathbf{z})\} = \mathbf{z}'\gamma_k$. Then the vector of parameters γ_k can be estimated by OLS – one OLS equation is estimated for each $a_i = k, k \neq HH$.

The probability function (A.5.1) and the Markovian transitions for actions and shares are used to simulate moments implied by the model. Starting from the initial state vector for each firm in each supermarket we forward simulate 1000 paths of 200 periods of actions and shares and computed profits for each period by averaging period profits for each path.

A.6 Additional tables

Table A.6.1: Annual expenditure shares in the 500g spreadable segment.

PRODUCTS BY MANUFACTURER	Year				2009-2012
	2009	2010	2011	2012	
ASDA					
Asda Store Brand	10.0%	9.3%	6.4%	5.8%	7.7%
Arla	35.6%	42.3%	39.5%	50.6%	42.1%
ANCHOR SPRDBL 500GM	10.7%	11.0%	10.8%	13.9%	11.6%
LURPAK SPRDBL DANISH 500GM	24.9%	31.3%	28.7%	36.7%	30.5%
Dairy Crest	21.0%	14.9%	16.6%	21.3%	18.3%
CLOVER DAIRY SPREAD 500GM	6.4%	6.6%	7.2%	11.3%	7.9%
COUNTRY LIFE SPRDBL 500GM	14.6%	8.3%	9.4%	10.0%	10.4%
Unilever	33.4%	33.5%	37.5%	22.3%	31.8%
FLORA 500GM	16.4%	14.1%	15.2%	21.6%	16.8%
I.C.B.I.N.B SPRD 500GM	17.1%	19.5%	22.2%	0.7%	15.0%
MORRISONS					
Morrisons Store Brand	9.9%	11.3%	9.0%	5.9%	9.0%
Arla	35.2%	36.2%	35.8%	37.0%	36.0%
ANCHOR SPRDBL 500GM	9.2%	8.8%	8.4%	10.1%	9.1%
LURPAK SPRDBL DANISH 500GM	26.0%	27.4%	27.3%	26.9%	26.9%
Dairy Crest	20.2%	18.3%	21.8%	23.3%	21.0%
CLOVER DAIRY SPREAD 500GM	12.0%	13.5%	14.2%	14.8%	13.7%
COUNTRY LIFE SPRDBL 500GM	8.2%	4.8%	7.5%	8.6%	7.3%
Unilever	34.7%	34.2%	33.4%	33.7%	34.0%
FLORA 500GM	26.9%	22.4%	21.6%	24.6%	23.8%
I.C.B.I.N.B SPRD 500GM	7.8%	11.8%	11.8%	9.1%	10.2%
SAINSBURY'S					
Sainsbury's Store Brand	15.2%	16.4%	17.4%	14.9%	16.1%
Arla	35.3%	38.3%	37.3%	40.3%	37.9%
ANCHOR SPRDBL 500GM	13.5%	14.3%	12.8%	14.1%	13.6%
LURPAK SPRDBL DANISH 500GM	21.9%	24.0%	24.5%	26.3%	24.3%
Dairy Crest	17.5%	15.6%	19.5%	19.2%	18.0%
CLOVER DAIRY SPREAD 500GM	8.1%	9.3%	11.2%	11.4%	10.1%
COUNTRY LIFE SPRDBL 500GM	9.3%	6.3%	8.3%	7.8%	7.9%
Unilever	32.0%	29.6%	25.8%	25.6%	28.0%
FLORA 500GM	21.5%	19.5%	15.9%	17.7%	18.5%
I.C.B.I.N.B SPRD 500GM	10.5%	10.1%	9.9%	7.8%	9.5%
TESCO					
Tesco Store Brand	17.7%	12.9%	15.5%	13.5%	14.9%
Arla	33.8%	40.2%	40.0%	39.0%	38.4%
ANCHOR SPRDBL 500GM	10.3%	13.7%	12.7%	11.2%	12.0%
LURPAK SPRDBL DANISH 500GM	23.6%	26.5%	27.3%	27.8%	26.4%
Dairy Crest	16.1%	16.7%	16.6%	17.2%	16.6%
CLOVER DAIRY SPREAD 500GM	9.1%	10.0%	9.3%	10.6%	9.7%
COUNTRY LIFE SPRDBL 500GM	7.0%	6.7%	7.3%	6.6%	6.9%
Unilever	32.4%	30.2%	28.0%	30.4%	30.1%
FLORA 500GM	24.3%	19.8%	16.1%	22.2%	20.4%
I.C.B.I.N.B SPRD 500GM	8.0%	10.4%	11.9%	8.1%	9.7%


Note: Calculations based on a subsample of products used to estimate the dynamic game. Source: own calculation using Kantar Worldpanel data.

Table A.6.2: Price levels.

PRODUCTS BY MANUFACTURER	MEANS		MEDIAN		MIN/MAX	
	<i>PH</i>	<i>PL</i>	<i>PH</i>	<i>PL</i>	<i>PH</i>	<i>PL</i>
ASDA						
Asda Store Brand	1.02		1.00			
Arla						
ANCHOR	2.51	1.82	2.60	2.00	2.90	1.00
LURPAK	2.63	2.10	2.58	2.00	2.98	1.50
Dairy Crest						
CLOVER	1.73	1.30	1.75	1.38	2.00	1.00
COUNTRY LIFE	2.42	1.85	2.39	2.00	2.68	1.00
Unilever						
FLORA	1.40	1.00	1.38	1.00	1.70	0.83
ICBINB	1.22	1.09	1.24	1.00	1.45	0.50
MORRISONS						
Morrisons Store Brand	1.09		1.08			
Arla						
ANCHOR	2.55	1.92	2.60	2.00	2.90	1.50
LURPAK	2.71	2.11	2.80	2.00	3.00	1.50
Dairy Crest						
CLOVER	1.75	1.15	1.75	1.00	2.00	0.70
COUNTRY LIFE	2.45	1.83	2.39	2.00	2.85	1.10
Unilever						
FLORA	1.47	0.94	1.40	1.00	1.70	0.70
ICBINB	1.21	0.82	1.25	1.00	1.45	0.50
SAINSBURY'S						
Sainsbury's Store Brand	1.13		1.10			
Arla						
ANCHOR	2.58	2.03	2.60	2.00	3.00	1.50
LURPAK	2.71	2.17	2.80	2.00	3.00	1.50
Dairy Crest						
CLOVER	1.75	1.22	1.75	1.00	2.00	0.85
COUNTRY LIFE	2.47	1.89	2.48	2.00	2.85	1.00
Unilever						
FLORA	1.48	0.96	1.49	1.00	1.70	0.75
ICBINB	1.27	0.79	1.25	1.00	1.80	0.54
TESCO						
Tesco Store Brand	1.02		1.00			
Arla						
ANCHOR	2.59	1.84	2.60	2.00	2.90	1.00
LURPAK	2.73	1.95	2.80	2.00	2.98	1.40
Dairy Crest						
CLOVER	1.74	1.18	1.75	1.00	2.00	0.75
COUNTRY LIFE	2.42	1.76	2.40	2.00	2.85	1.10
Unilever						
FLORA	1.49	1.01	1.46	1.00	1.70	0.75
ICBINB	1.24	0.88	1.24	1.00	1.80	0.54

Note: All prices given in GBP. First four columns show prices calculated as 200-week averages/medians conditional on promotional status. For store brand there are no price promotions, so it is an unconditional mean/median. Prices in the last two columns are calculated as highest/lowest price observed in the sample period conditional on sale/no sale.

Figure A.6.1: Histograms with the number of products on sale by firm.



Note: Figure constructed using the universe of all 500g spreadable products by recording the promotional flags for each of the products. E.g. for Arla there were 19 weeks with no product on sale, 84 weeks with 1 brand on sale, 68 weeks with 2 brands on sale etc. If the numbers do not sum to 200 for certain manufacturers it is an indication that we did not observe any purchases their brands in the data in all weeks.

A.7 Additional results

Table A.7.1: Demand estimates (alternative definition of loyalty).

	ASDA	MORRISONS	SAINSBURY'S	TESCO
δ_{Anchor}	-2.628 [-2.753; -2.503]	-2.387 [-2.541; -2.233]	-3.143 [-3.290; -2.995]	-3.600 [-3.676; -3.523]
δ_{Lurpak}	-2.097 [-2.220; -1.975]	-1.744 [-1.893; -1.595]	-2.903 [-3.054; -2.752]	-3.260 [-2.49; -2.35]
δ_{Clover}	-3.223 [-3.323; -3.123]	-2.730 [-2.834; -2.626]	-3.785 [-3.892; -3.677]	-3.879 [-3.936; -3.822]
$\delta_{Country\ Life}$	-2.865 [-2.984; -2.745]	-2.794 [-2.938; -2.650]	-3.590 [-3.736; -3.443]	-4.156 [-4.234; -4.077]
δ_{Flora}	-3.055 [-3.138; -2.973]	-2.631 [-2.720; -2.542]	-3.455 [-3.550; -3.360]	-3.620 [-3.669; -3.570]
δ_{ICBINB}	-3.138 [-3.211; -3.065]	-3.036 [-3.117; -2.955]	-3.960 [-4.050; -3.870]	-3.921 [-3.969; -3.874]
δ_{SB}	-2.301 [-2.366; -2.237]	-2.280 [-2.350; -2.210]	-2.021 [-2.092; -1.950]	-3.260 [-3.333; -3.186]
η	-0.966 [-1.019; -0.912]	-0.905 [-0.964; -0.846]	-0.587 [-0.646; -0.528]	-0.409 [-0.440; -0.377]
γ	2.564 [2.527; 2.601]	2.086 [2.042; 2.129]	2.779 [2.741; 2.818]	2.569 [2.546; 2.592]
N	104,946	71,294	102,939	280,828
pseudo- R^2	0.285	0.329	0.152	0.133

Note: Estimates obtained using the alternative definition of loyalty (see appendix A). All parameters are significantly different from 0 at the 1% level. 95% confidence intervals reported below estimated coefficients, constructed using robust standard errors. *SB* denotes store brand.

Table A.7.2: Demand estimates (heterogenous γ).

	ASDA	MORRISONS	SAINSBURY'S	TESCO
δ_{Anchor}	-2.975 [-3.109; -2.841]	-3.042 [-3.210; -2.875]	-3.575 [-3.727; -3.424]	-4.100 [-4.179; -4.020]
δ_{Lurpak}	-2.202 [-2.327; -2.078]	-2.176 [-2.332; -2.020]	-3.177 [-3.329; -3.026]	-3.545 [-3.618; -3.471]
δ_{Clover}	-3.148 [-3.253; -3.043]	-2.852 [-2.962; -2.741]	-3.821 [-3.931; -3.711]	-4.085 [-4.145; -4.025]
$\delta_{Country\ Life}$	-3.898 [-4.027; -3.770]	-3.348 [-3.505; -3.190]	-4.322 [-4.484; -4.160]	-4.932 [-5.020; -4.844]
δ_{Flora}	-2.299 [-2.376; -2.222]	-2.117 [-2.197; -2.036]	-2.763 [-2.848; -2.679]	-2.925 [-2.969; -2.882]
δ_{ICBINB}	-2.501 [-2.568; -2.434]	-2.840 [-2.920; -2.759]	-3.447 [-3.532; -3.362]	-3.701 [-3.746; -3.656]
δ_{SB}	-2.815 [-2.885; -2.746]	-2.997 [-3.075; -2.918]	-2.576 [-2.650; -2.503]	-3.051 [-3.086; -3.016]
η	-0.744 [-0.800; -0.688]	-0.656 [-0.718; -0.594]	-0.295 [-0.355; -0.235]	-0.130 [-0.162; -0.098]
γ_{Anchor}	4.006 [3.870; 4.143]	3.886 [3.712; 4.061]	4.012 [3.904; 4.121]	4.095 [4.024; 4.165]
γ_{Lurpak}	3.547 [3.459; 3.635]	3.382 [3.275; 3.489]	3.598 [3.506; 3.690]	3.658 [3.606; 3.711]
γ_{Clover}	3.384 [3.258; 3.511]	3.402 [3.289; 3.515]	3.917 [3.805; 4.030]	4.146 [4.081; 4.211]
$\gamma_{CountryLife}$	3.988 [3.850; 4.126]	3.846 [3.645; 4.047]	4.768 [4.614; 4.921]	4.948 [4.844; 5.053]
γ_{Flora}	2.445 [2.373; 2.518]	2.273 [2.202; 2.345]	2.673 [2.607; 2.739]	2.666 [2.628; 2.705]
γ_{ICBINB}	2.981 [2.917; 3.046]	3.087 [2.995; 3.179]	2.977 [2.881; 3.073]	3.458 [3.402; 3.513]
γ_{SB}	2.857 [2.818; 2.897]	3.302 [3.201; 3.404]	1.851 [1.782; 1.921]	2.857 [2.818; 2.897]
N	104,946	71,294	102,939	280,828
pseudo- R^2	0.289	0.367	0.150	0.187

Note: All parameters are significantly different from 0 at the 1% level. 95% confidence intervals reported below estimated coefficients, constructed using robust standard errors. *SB* denotes store brand.

Table A.7.3: Multinomial logit CCP estimates.

	Arla			Dairy Crest			Unilever		
	HL	LH	LL	HL	LH	LL	HL	LH	LL
a_{t-1}									
Arla: HL	2.064*** (0.08)	0.592* (0.32)	2.091*** (0.35)	-0.679 (0.61)	0.116 (0.24)	0.228 (0.46)	-0.001 (0.14)	0.333 (0.71)	-0.051 (0.73)
Arla: LH	-0.032 (0.32)	2.385*** (0.29)	2.450*** (0.46)	-0.398 (0.54)	-0.452 (0.37)	-0.466 (0.45)	-0.232 (0.37)	0.495 (0.78)	-0.070 (0.95)
Arla: LL	2.869*** (0.83)	3.018*** (0.65)	5.031*** (0.79)	0.124 (0.46)	-0.728 (0.56)	-0.219 (0.70)	-0.148 (0.32)	0.635 (0.64)	-0.059 (0.58)
DC: HL	0.633 (0.49)	-0.107 (0.25)	-0.308 (0.35)	3.283*** (0.34)	0.805 (0.56)	2.668*** (0.32)	0.120 (0.13)	-0.005 (0.46)	-0.620 (0.40)
DC: LH	0.133 (0.28)	-0.403 (0.31)	-0.087 (0.30)	0.846 (0.55)	2.732*** (0.45)	2.074*** (0.30)	0.146 (0.21)	-0.339 (0.54)	-0.720* (0.40)
DC: LL	-0.205 (0.20)	-0.569 (0.41)	-1.062** (0.52)	2.312*** (0.53)	2.780*** (0.34)	4.387*** (0.61)	0.035 (0.20)	-0.221 (0.45)	-0.999 (0.69)
Unilever: HL	-0.082 (0.27)	0.129 (0.15)	0.323* (0.19)	-0.072 (0.35)	0.340 (0.38)	-0.948*** (0.35)	2.512*** (0.28)	-0.059 (0.87)	1.752*** (0.28)
Unilever: LH	-0.313 (0.58)	-0.068 (0.28)	0.474* (0.27)	0.087 (0.26)	-0.379 (0.39)	-0.404 (0.40)	0.583 (0.40)	3.023*** (0.15)	3.037*** (0.23)
Unilever: LL	-0.812* (0.48)	-0.122 (0.11)	-0.055 (0.59)	-0.548* (0.32)	-0.077 (0.42)	-0.613 (0.51)	2.034* (1.14)	1.487** (0.59)	4.261*** (0.83)
s_{t-1}									
ANCHOR	46.893** (18.32)	40.588 (32.63)	58.726 (37.59)	14.085 (25.40)	-14.866 (17.11)	-20.447** (9.74)	-5.319 (19.72)	0.058 (24.10)	44.926* (26.81)
LURPAK	39.537*** (14.93)	19.496 (13.72)	19.526** (8.34)	2.885 (11.76)	33.039*** (6.75)	19.656* (10.25)	12.349 (19.99)	14.877* (8.21)	33.932*** (10.48)
CLOVER	-15.452* (8.95)	5.741 (5.75)	10.781 (6.90)	8.741** (4.28)	-12.049*** (4.59)	12.354* (6.58)	-5.322 (6.03)	4.464 (3.40)	-4.045 (6.78)
COUNTRY LIFE	-25.289*** (7.41)	6.071 (11.71)	-2.554 (17.65)	28.405* (16.15)	28.989 (21.78)	42.215* (22.18)	-57.456*** (6.57)	5.898 (15.09)	42.557** (20.47)
FLORA	3.161 (4.54)	3.139 (5.75)	2.408 (3.84)	-3.202 (6.58)	-13.367*** (4.69)	2.528 (8.63)	6.453 (5.80)	3.976 (9.42)	9.420* (5.67)
ICBINB	0.305 (5.44)	-4.137* (2.42)	-1.185 (2.91)	-3.324 (3.75)	3.058 (6.31)	3.853 (2.96)	4.523 (5.63)	12.564*** (1.65)	13.773*** (2.73)
STORE BRAND	-1.132 (6.05)	-2.270 (8.26)	-18.353 (14.07)	-6.775 (15.52)	-4.959 (4.18)	1.167 (10.09)	-2.047 (5.64)	-22.016*** (7.62)	-25.182*** (8.39)
MORRISONS	-0.181 (0.13)	-0.519*** (0.17)	-0.457** (0.22)	0.589*** (0.20)	-0.094 (0.17)	0.122 (0.22)	0.723*** (0.11)	-0.579*** (0.09)	-0.001 (0.13)
SAINSBURY'S	-0.289* (0.17)	-0.905** (0.38)	-1.754*** (0.61)	0.223 (0.45)	-0.607** (0.29)	-1.109*** (0.33)	0.707** (0.29)	-0.574 (0.48)	-1.828*** (0.52)
TESCO	1.135*** (0.18)	0.532 (0.37)	0.921** (0.42)	0.509 (0.42)	-0.249* (0.13)	0.184 (0.34)	1.086* (0.20)	0.109 (0.34)	0.687*** (0.20)
Constant	-2.042*** (0.56)	-1.223** (0.62)	-3.031*** (0.56)	-2.114*** (0.80)	-1.155*** (0.41)	-3.488*** (0.41)	-2.303*** (0.37)	-1.712*** (0.43)	-4.349*** (1.09)

Note: For all 3 players (Arla, Dairy Crest, Unilever) *HH* is the reference action. *H* stands for high and *L* low price, for the two products each firm is selling. Arla: Anchor and Lurpak, Dairy Crest: Clover and Country Life, Unilever: Flora and I Can't Believe It's Not Butter (ICBINB). Last panel of the table shows supermarket fixed effects to reflect the fact that different equilibrium strategies can be played in different markets. Asda is the reference market there. $N = 703$. Significance levels: *** 1%, ** 5%, * 10%.

Figure A.7.1: Actions played by firms: model vs. data.

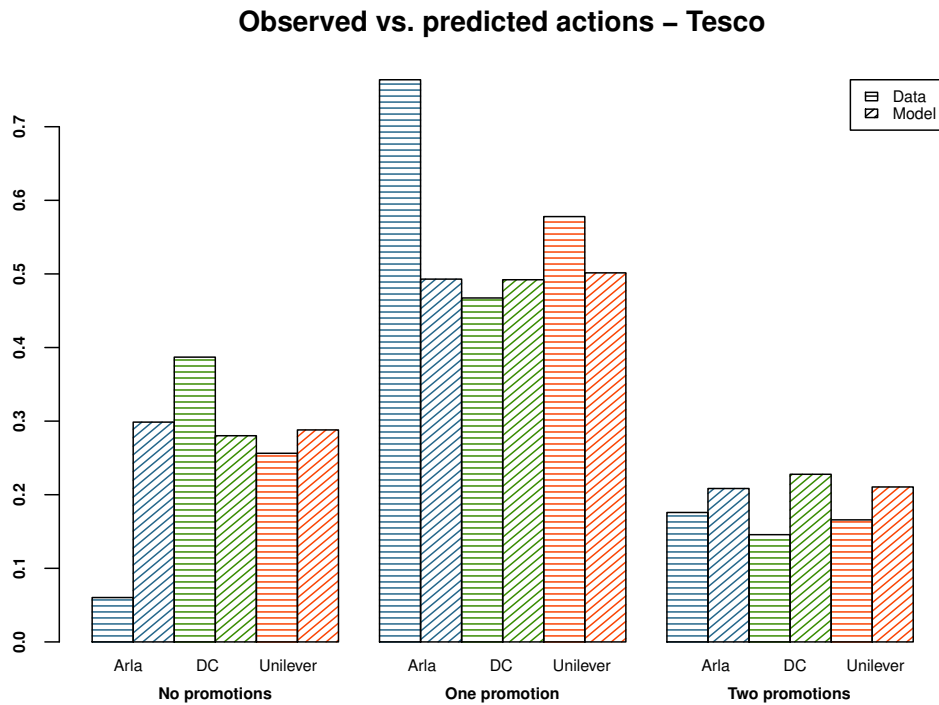
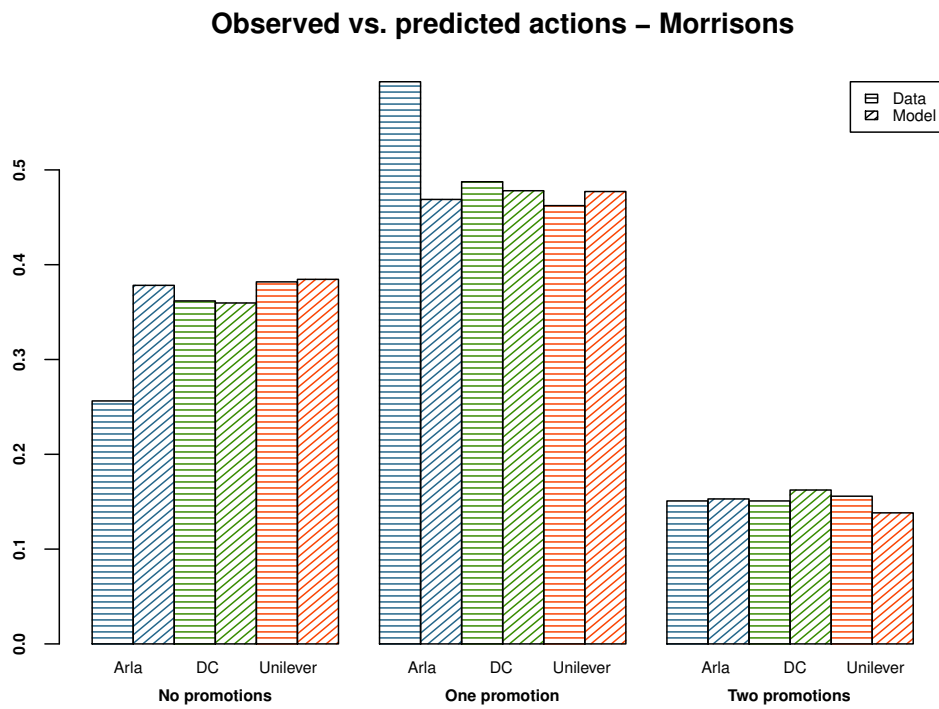
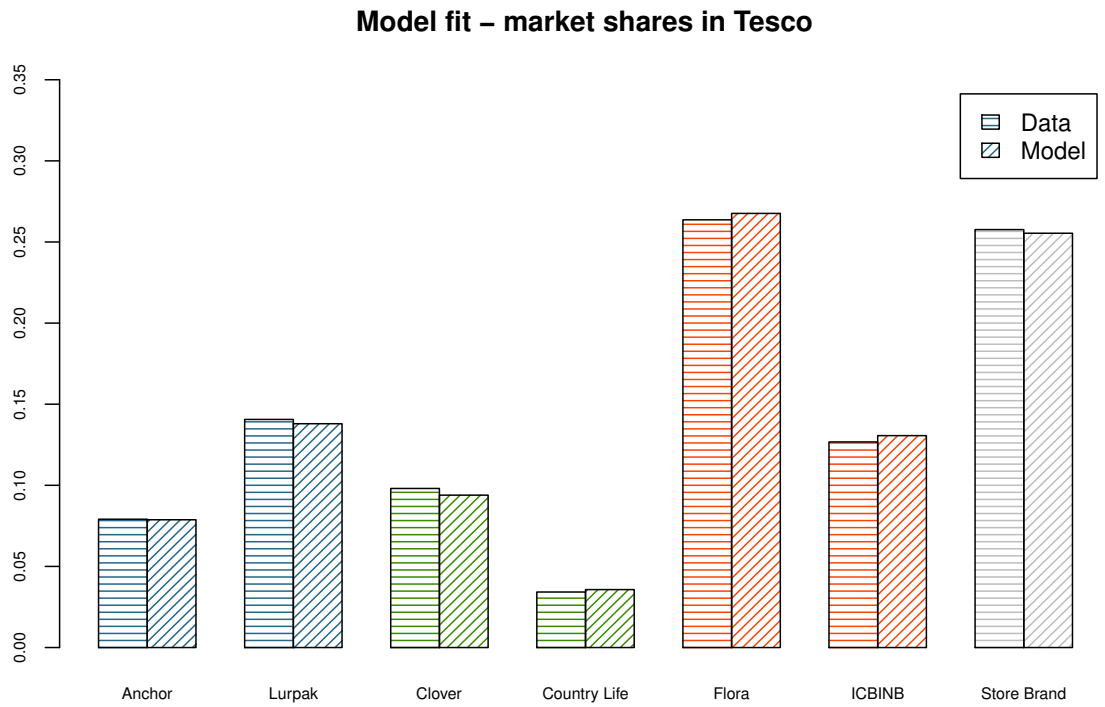
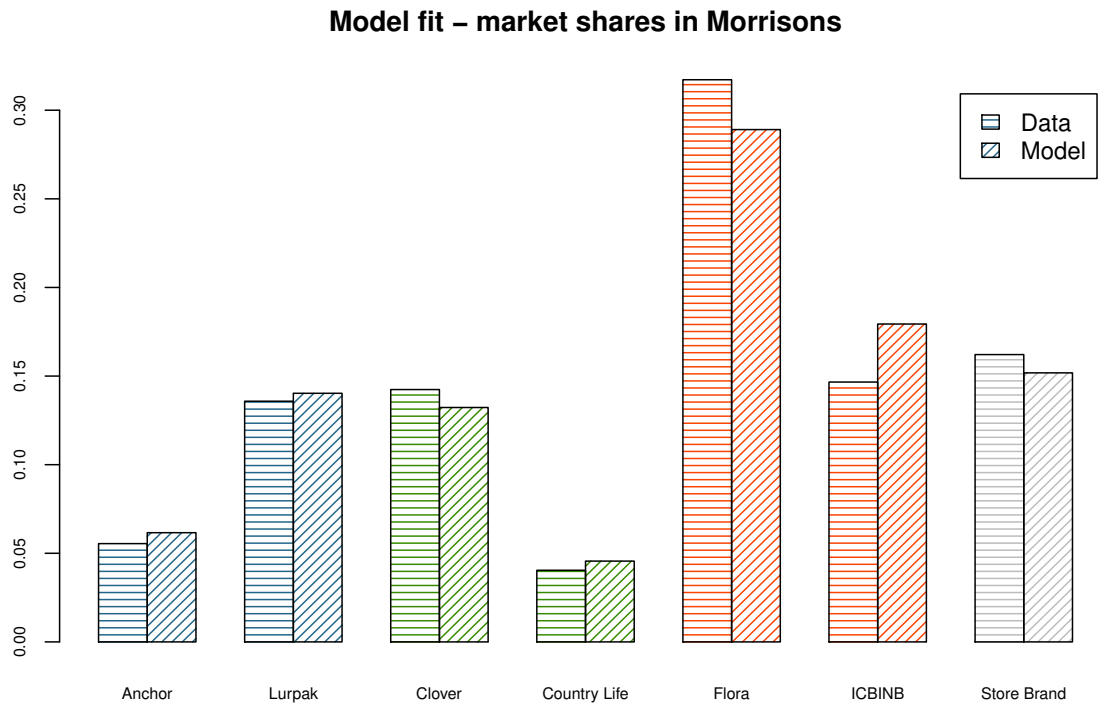


Figure A.7.2: Market shares by brand: model vs. data.



Appendix B

Supplementary material for chapter 2

B.1 Proofs

The equivalent of Proposition 1 of GPV (2000)

The goal of this proposition is to study the smoothness of the density of prices $f(\cdot)$, implied by some assumptions on the latent density of marginal costs, $h(\cdot)$. This is needed to study the uniform convergence rate of the nonparametric estimator. For simplicity, we assume that the number of firms is fixed at K . Paraphrasing assumptions A1-A2 of GPV, we have:

Assumptions:

A1 (IPV): r_i 's are independently and identically distributed as $\mathcal{H}(\cdot)$ with density $h(\cdot)$,

A2 (Finite support): For all K , $\text{supp}(\mathcal{H}) = [\underline{R}, \bar{R}]$ is a compact subset of \mathbb{R}_+ .

A3 (Non-zero density): For all $r \in \text{supp}(\mathcal{H})$, $h(r) \geq c_h > 0$

A4 (Smoothness): $\mathcal{H}(\cdot)$ admits up to $R + 1$ continuous bounded partial derivatives on $\text{supp}(\mathcal{H})$, with $R \geq 1$.

Assumptions A1-A4 allow us to formulate the following proposition:

Proposition 1:

The distribution of prices $\mathcal{F}(\cdot)$ satisfies:

(i) $\text{supp}(\mathcal{F}) = \{p : p \in [\underline{P}, \bar{P}]\}$ with $\bar{P} > \underline{P}$ and $\bar{P} = \bar{R}$.

(ii) for all $p \in \text{supp}(\mathcal{F})$, $f(p) \geq c_f > 0$

(iii) $\mathcal{F}(\cdot)$ admits up to $R + 1$ continuous bounded partial derivatives on $\text{supp}(\mathcal{F})$

(iv) $f(\cdot)$ admits up to $R + 1$ continuous bounded partial derivatives on every closed subset of the interior of $\text{supp}(\mathcal{F})$

Proof. Monotonicity of the optimal pricing strategy and the boundary condition $\beta(\bar{R}) = \bar{R}$ make (i) trivially satisfied. To prove (ii), note that $f(p) = \frac{h(\xi(p))}{\beta'(\xi(p))}$. h by A3 is bounded away from 0 and by A4 so the derivative of the optimal pricing strategy $\beta'(\cdot)$, as long as $q_1 \neq 1$. Since we focus on equilibria with price dispersion, $\beta'(\cdot) > 0$ holds almost everywhere. Hence $f(p)$ is also bounded away from 0. To prove (iii) it is enough to observe that $\mathcal{F}(p) = \mathcal{H}(\xi(p))$, where both \mathcal{H} and ξ have $R + 1$ continuous bounded derivatives. Finally, since for any p in the interior of $\text{supp}(\mathcal{F})$, $(p - \xi(p)) > 0$, the expression for the inverse of the optimal pricing strategy can be rearranged to yield the density:

$$f(p) = \frac{\sum_{k=1}^K q_k k (1 - \mathcal{F}(p))^{k-1}}{(p - \xi(p)) \sum_{k=2}^K q_k k (k-1) (1 - \mathcal{F}(p))^{k-2}} \quad [\text{B.1.1}]$$

Since $\xi(\cdot) = \beta^{-1}(\cdot)$ and $\beta(\cdot)$ has the same smoothness as \mathcal{H} , then $\xi(\cdot)$ is $R + 1$ -times continuously differentiable. But since according to (iii), $\mathcal{F}(\cdot)$ is also $R + 1$ -times differentiable, then so is $f(\cdot)$ and (iv) follows.

Appendix C

Supplementary material for chapter 3

C.1 Data and reduced form results

C.1.1 Data and summary statistics

Table C.1.1 summarises our main variables of interest by broker usage for different types of mortgage products– the two-, three-, and five-year FRMs. There is variation in loan size, fees, and offered interest rate across product type, but between borrowers who go direct or use brokers. The last row in the table shows monthly interest payments normalised by the size of the loan, which is our preferred measure of calculating mortgage cost. Section 3.2.3 outlines the calculation in detail.

Table C.1.1: Summary statistics

	2 yr FRM		3 yr FRM		5 yr FRM	
	Direct	Broker	Direct	Broker	Direct	Broker
Interest (bps)	232	213	245	235	259	247
Loan (£)	161,070	185,578	166,804	176,525	145,680	167,747
Loan Fees (£)	373.97	575.18	434.58	578.43	495.12	620.56
Monthly payment (£)	788.47	781.00	831.52	796.33	782.82	791.66
Monthly interest (£)	317.85	358.93	345.73	363.63	316.03	355.56
Normalised payment (£)	305.84	298.95	314.16	312.54	331.70	320.59

Note: Interest is the interest rate in basis points. Loan is the size of the mortgage issued by the bank. Monthly payment is the payment of capital and interest during the initial contract period of the loan, excluding lender fees. Monthly interest is the component of the monthly payment that goes towards payment of the interest, and includes the fees. Normalised interest payment is the monthly interest payment normalised to take into account the size of the loan.

Mortgage contracts in the UK are short-term, with an initial duration of 2-, 3-, or 5-years. Following the expiration of the initial period, and if the household does not refinance, the mortgage contract reverts to the bank’s posted rate, or Standard Variable Rate (SVR). There are two types of contracts in the UK: fixed and variable. Fixed rate mortgages (FRM) have a fixed interest

rate during the initial period, while adjustable rate mortgages (ARM) have a fluctuating rate that is a discount off of the SVR. Mortgage rates are arranged according to the length of the initial period and by LTV band. The longer the period and the higher the LTV, the more expensive the product. Table C.1.2 shows that, on average in our sample, households pay 230 basis points on their mortgage product, but that there is a spread of 280 basis points between the 2-year FRM at 70% LTV (cheapest) and the 5-year FRM at 95% LTV. Given that yield curves were roughly flat during this period, spreads across products have remained more or less constant.

Table C.1.2: Interest Rates by LTV and Rate Duration

	2 yr FRM	3 yr FRM	5 yr FRM	Total
≤ 70	1.8	2.2	2.3	2.0
71 - 75	1.8	2.2	2.5	2.0
76 - 80	1.9	2.4	2.6	2.1
81 - 85	2.1	2.5	2.8	2.2
86 - 90	2.8	3.0	3.3	2.9
91 - 95	3.0	4.0	4.6	4.0
Total	2.2	2.4	2.5	2.3

Just over one-third of our sample are FTB, with the remainder either moving home or remortgaging their current home. But there is variation in the distribution of mortgagors at different LTV bands. Table C.1.3 shows that 80% of mortgagors on 95% LTV products are FTB, whereas 80% of mortgagors who took out an LTV of 70% or less are non-FTB.

Different banks also specialise in different products, with the share of longer term products more likely to be offered by some banks over others. This can be seen in table C.1.4

Table C.1.3: Share by Household type and LTV

	Non-FTB	FTB	Total
≤ 70	82	18	100
71 - 75	60	40	100
76 - 80	68	32	100
81 - 85	56	44	100
86 - 90	36	64	100
91 - 95	19	81	100
Total	64	36	100

Table C.1.4: Share by Bank and Product Type

	2 yr FRM	3 yr FRM	5 yr FRM	Total
Bank 1	76.61	0.89	22.50	100
Bank 2	67.48	2.26	30.26	100
Bank 3	58.80	9.75	31.45	100
Bank 4	66.22	4.71	29.07	100
Bank 5	44.19	10.74	45.07	100
Bank 6	72.30	1.42	26.28	100
Total	66.27	4.33	29.40	100

C.1.2 Estimation sample

We restrict our sample to standard¹ fixed rate mortgage products with two-, three-, and five-year durations; and to loan sizes less than £1M. This leaves us with about 82% of the sample (1.7M loans) for analysis. We further restrict our sample to the six largest mortgage providers which made up about 75% (or 1.3M loans) in 2016 and 2017. The differences between the raw and final sample are tabulated in table C.1.5.

Table C.1.5: Raw and Final Sample

	Big Six	%	Raw Sample	%
Total	1,539,009	100.00	2,138,754	100.00
Interest-only mortgages	43,276	2.81	81,482	3.81
Non-FRM	114,099	7.41	152,856	7.15
Not 2, 3, 5yrs	61,765	4.01	141,054	6.60
£1M+ loan	4,186	0.27	5,886	0.28
Outliers	6,606	0.43	13,892	0.65
Final Sample	1,309,077	85.06	1,743,584	81.52

C.1.3 Probability of using a broker

Table C.1.6 reports the estimates from a linear probability model where we regressed the indicator whether the contract was brokered on a number of personal and product characteristics. The first observation is that the signs are in line with intuition. For example, lower income, first time buyers, the employed, and older mortgagors are more likely to use a broker. Moving to column 2, adding product characteristics shows that mortgagors who took longer term contracts were less likely to visit brokers (the causality may also be in the other direction, so we interpret the results in terms of conditional correlations, rather than causal relationships). In fact, a recent FCA investiga-

¹These are products that include repayment of the capital.


tion (Ischenko and Nieboer, 2018) hypothesises that brokers might be more likely to suggest 2-year contracts knowing that this makes borrowers use their services more frequently in the future. A longer mortgage term is also associated with increased probability of using brokers. However, column 2 also shows that when product characteristics are added, the sign on LTV indicators is reversed from positive (column 1) to negative. In fact, the higher the LTV the less likely a household uses a broker. This may be for a number of reasons, for example, households on low LTV products typically have smaller absolute loans, therefore the costs of visiting a broker and paying a lump-sum is relatively higher. Finally, column 3 shows that even after controlling for regional fixed effects, the coefficients remain unchanged and the R^2 remains low, so the observables are rather poor predictors for broker use.

Table C.1.6: Probability of using a broker

	(1)	(2)	(3)
Dependent var:	Personal	Product	Regional
Used a broker	Characteristics	Characteristics	Characteristics
Income	-0.002***	-0.016***	-0.037***
First Time Buyer	0.048***	0.016***	0.009***
Aged 25 - 29	0.024***	0.028***	0.024***
Aged 30-34	0.042***	0.068***	0.063***
Aged 35- 39	0.049***	0.121***	0.115***
Aged 40 - 45	0.036***	0.179***	0.172***
Aged 45+	-0.022***	0.251***	0.241***
71 - 75 LTV	0.131***	0.063***	0.078***
76 - 80 LTV	0.042***	-0.029***	-0.012***
81 - 85 LTV	0.075***	-0.019***	-0.000
86 - 90 LTV	0.035***	-0.066***	-0.041***
91 - 95 LTV	0.028***	-0.111***	-0.084***
Employed	-0.054***	-0.046***	-0.045***
Mortgage Term		0.021***	0.020***
3 Year FRM		-0.229***	-0.228***
5 Year FRM		-0.187***	-0.184***
Flexible Mortgage		0.086***	0.083***
Urban area			-0.011***
Regional FE	No	No	Yes
Observations	1,309,067	1,309,067	1,307,538
R^2	0.020	0.124	0.130

Note: *** denotes 1% significance level. Robust standard errors used.

Figure C.1.1: Distributions of predicted probabilities of using a broker.



Note: Density estimates of the distributions of $\widehat{\Pr}(d_i = \text{broker}|\mathbf{X})$ based on the LPM in the third column of table C.1.6 for the brokered and direct subsamples.

C.1.4 Robustness checks

This section presents robustness checks, which examine potential effects of procurement fees paid by the lenders to the brokers. The first two tables display the results of the regression of prices on brokered dummy (table 3.2 in the main text) for two subsamples of the data – C.1.7 only uses data on brokers who are not paid by the borrowers directly and are only compensated by the lenders, while C.1.8 only uses data on brokers who are not paid by the lenders and are only paid directly by the borrowers. The signs on the variables of interest are negative for all specifications and subsamples. This suggests that there is no evidence that different sources of compensation can alter brokers incentives to provide advice about cheaper products.

Table C.1.7: Price benefits of using a broker: brokers who do not charge the borrowers.

Dependent variable:	(1) Interest	(2) Interest	(3) Interest	(4) Monthly Payment	(5) Monthly Payment
Used a broker	-6.509*** (0.0902)	-6.370*** (0.0917)	-7.720*** (0.0927)	-2.091*** (0.143)	-2.210*** (0.153)
Lender Fees	Linear	Linear	Non-linear	-	-
Controls	Yes	Yes	Yes	Yes	Yes
Regional FE	No	Yes	Yes	No	Yes
Time FE	Yes	Yes	Yes	Yes	Yes
Observations	940,921	940,921	940,921	940,921	940,921
R^2	0.741	0.747	0.754	0.600	0.605

Note: *** denotes significant at 1% level. Robust standard errors in parentheses. Interest is measured in basis points. Monthly interest is the component of the initial monthly payment that goes towards payment of the interest, including lender fees, and normalised by the size of the loan. Controls are income, house price, loan size, LTV, first time buyer and mortgage term. Time fixed effects are at the monthly level. Regional fixed effects are at the Government Office Region level and include a flag for an urban region. Non-linearities in lender fees are controlled for using a fifth-order spline.

Table C.1.8: Price benefits of using a broker: brokers who are not paid by lenders.

Dependent variable:	(1) Interest	(2) Interest	(3) Interest	(4) Monthly Payment	(5) Monthly Payment
Used a broker	-6.181*** (0.180)	-2.443*** (0.177)	-4.390*** (0.182)	-6.151*** (0.244)	-1.046*** (0.245)
Lender Fees	Linear	Linear	Non-linear	-	-
Controls	Yes	Yes	Yes	Yes	Yes
Regional FE	No	Yes	Yes	No	Yes
Time FE	Yes	Yes	Yes	Yes	Yes
Observations	464,012	464,012	464,012	464,012	464,012
R^2	0.689	0.698	0.704	0.628	0.636

Note: *** denotes significant at 1% level. Robust standard errors in parentheses. Interest is measured in basis points. Monthly interest is the component of the initial monthly payment that goes towards payment of the interest, including lender fees, and normalised by the size of the loan. Controls are income, house price, loan size, LTV, first time buyer and mortgage term. Time fixed effects are at the monthly level. Regional fixed effects are at the Government Office Region level and include a flag for an urban region. Non-linearities in lender fees are controlled for using a fifth-order spline.

Tables C.1.9 and C.1.10 display alternative calculations of the value of information under the assumption that removing brokers would reduce lenders' costs by the expected amount of procurement fees. To implement this, we adjust each estimated $\mathcal{H}(\cdot|\mathbf{x}^H)$ by $\overline{\Delta^B} \cdot \overline{\phi}(\mathbf{x}^H)$, where the

first term is the average proportion of brokered loans with characteristics \mathbf{x}^H and the second term is the average observed procurement fee for a mortgage characterised by \mathbf{x}^H taken from the data. Since the result is a leftward shift of the entire distribution, equilibrium search behaviour does not change because proportions of borrowers searching different number of lenders depend on the differences in the expected prices and not the level of prices itself. The results of this exercise are valid under the assumption that the mortgage sold through a broker and directly is indeed the same product so any additional cost, such as the procurement fee if it is sold through a broker, is also indirectly passed onto consumers who obtain it directly from the lender.

The numbers in the tables below should be compared to tables 3.7 and 3.8 in the main text. Overall, adjusting for procurement fees reduces the value of information by about £15 through a smaller increase in prices (29% vs. 33%).

Table C.1.9: Value of information with adjusted costs: breakdown by borrower types.

	VOI	$\% \Delta p$	$\% \Delta SE$
Overall	97.38	+29.15%	+16.33%
Age			
<30	183.46	+51.18%	+65.49%
30+	78.76	+24.39%	+5.70%
Income			
Low	135.64	+36.53%	-18.62%
High	83.44	+26.46%	+50.43%
FTB			
FTB	111.57	+32.88%	+62.98%
Non-FTB	83.42	+25.51%	-0.67%
Location			
Urban	99.57	+29.42%	+26.05%
Rural	89.33	+28.18%	-19.46%

Note: Second column of the table reports the estimated average value of information as defined in equation (3.13) in GBP per month. The third and fourth columns report the average percentage change in prices and search expenditures, respectively. Calculations made by simulating new prices and search behaviour from the new equilibrium, assuming that lenders drew had the same cost draws as in the baseline scenario. Marginal cost distributions are adjusted to account for the fact that in a world without brokers, lenders do not pay procurement fees.

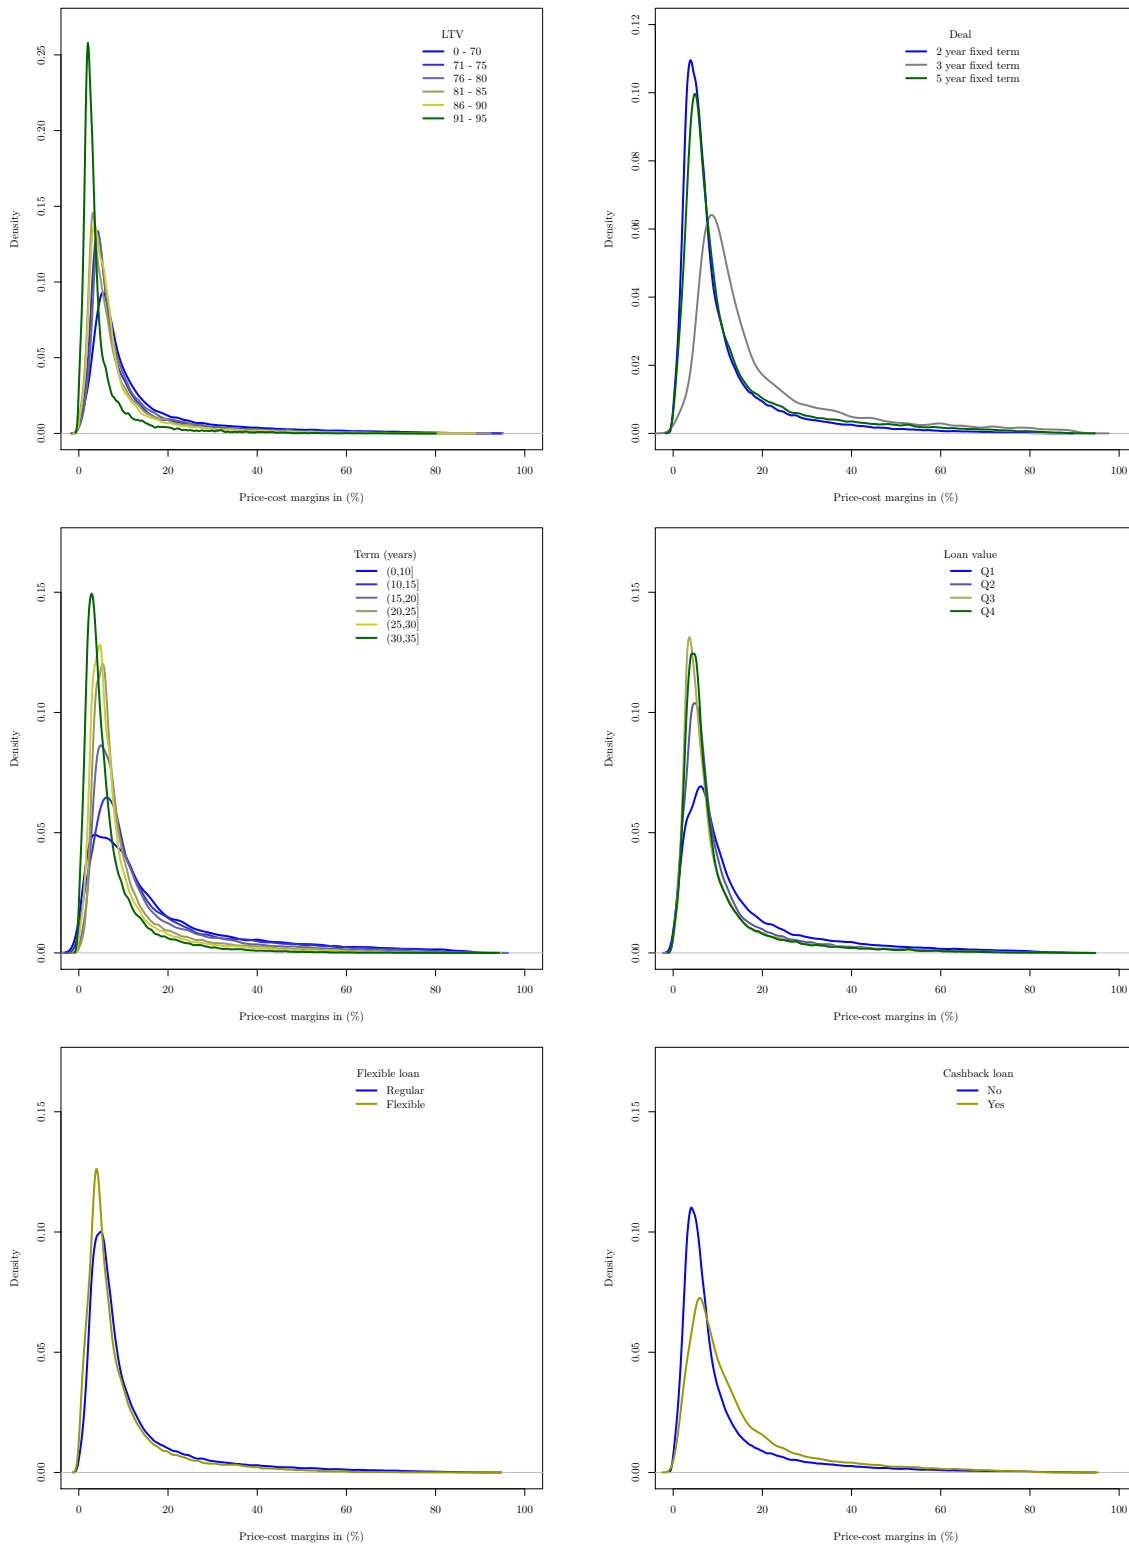
Table C.1.10: Value of information with adjusted costs: breakdown by loan characteristics.

	VOI	% Δp	% ΔSE
Overall	97.37	+29.15%	+16.33%
LTV			
≤ 70	69.65	+23.09%	-26.37%
71-75	132.21	+42.63%	+90.52%
76-80	51.31	+16.38%	+26.06%
81-85	163.04	+49.09%	+98.04%
86-90	151.47	+34.97%	+48.36%
91-95	45.49	+9.13%	-50.29%
Deal			
2 years	146.15	+43.38%	+48.25%
3 years	-9.45	-0.96%	-67.62%
5 years	-9.27	-2.05%	-52.44%
Term			
≤ 10 years	-2.69	+2.26%	-86.19%
(10;15]	-24.38	-5.34%	-77.46%
(15;20]	-9.36	-1.33%	-69.37%
(20;25]	60.89	+19.53%	-43.84%
(25;30]	108.85	+31.75%	+7.88%
(30;35]	342.77	+98.11%	+284.55%
Value			
Q1	101.83	+30.98%	-50.97%
Q2	66.17	+19.43%	-28.31%
Q3	103.31	+29.87%	+30.91%
Q4	116.71	+35.83%	+112.17%
Flexible			
Flexible	4.47	+4.13%	-82.27%
Regular	111.71	+33.01%	+31.69%
Cashback			
No cashback	106.57	+31.59%	-70.04%
Cashback	35.87	+12.84%	+29.33%

Note: Second column of the table reports the estimated average value of information as defined in equation (3.13) in GBP per month. The third and fourth columns report the average percentage change in prices and search expenditures, respectively. Calculations made by simulating new prices and search behaviour from the new equilibrium, assuming that lenders drew had the same cost draws as in the baseline scenario. Marginal cost distributions are adjusted to account for the fact that in a world without brokers, lenders do not pay procurement fees.

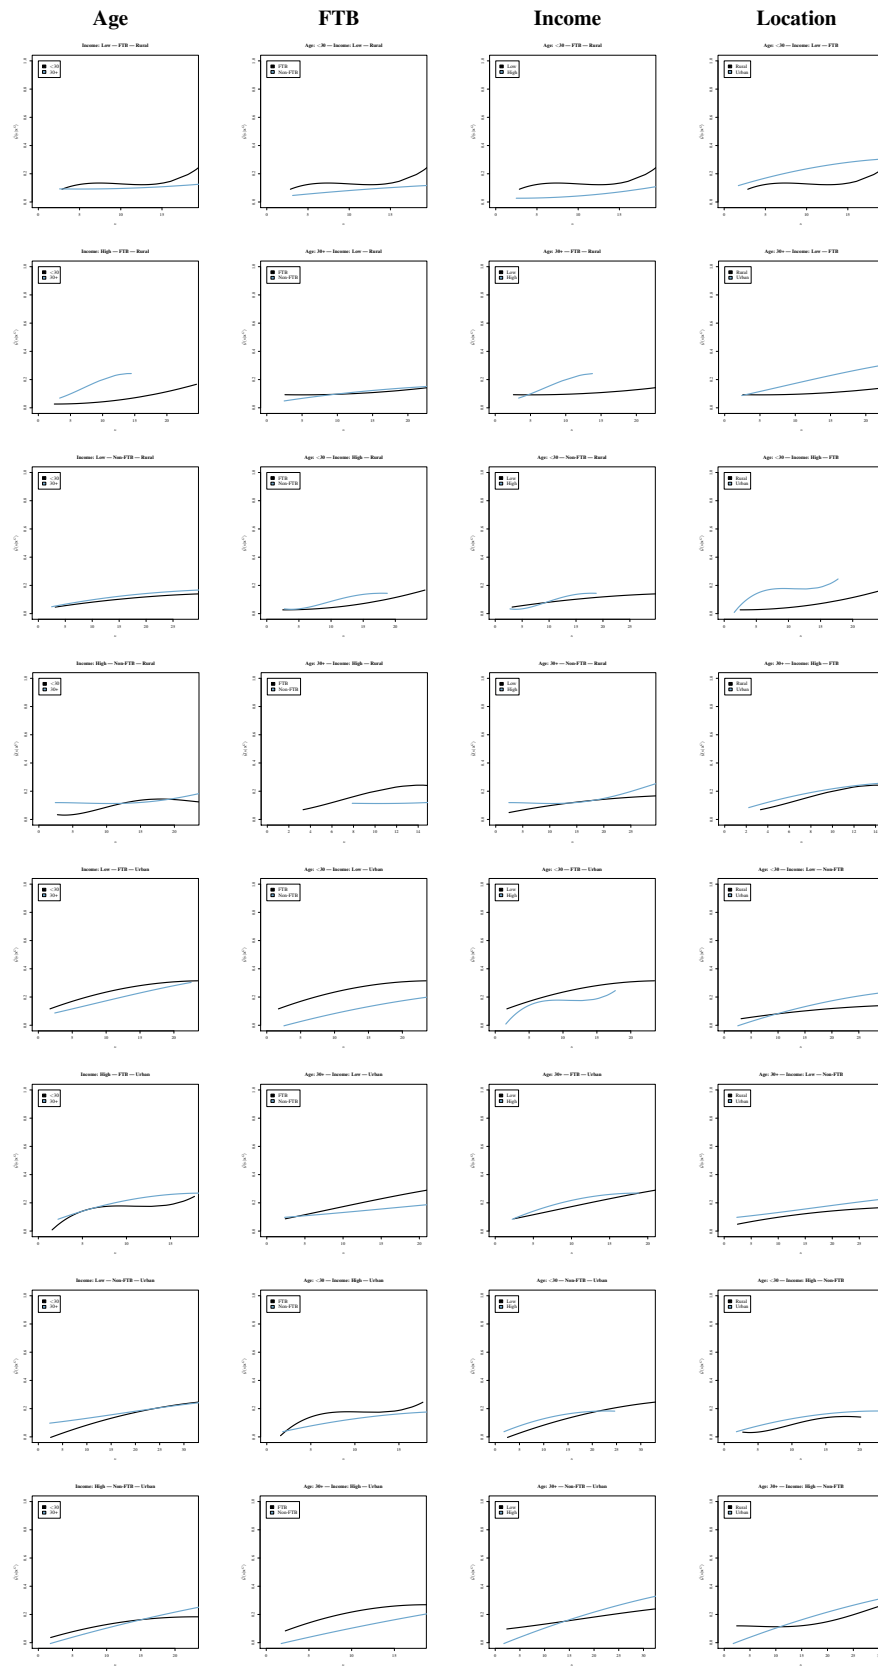
C.2 Additional figures

Figure C.2.1: Distributions of price-cost margins.



Note: Kernel estimate of the density of price-cost margins defined as $PCM_{ij} = \frac{p_{ij} - c_{ij}}{p_{ij}}$.

Figure C.2.2: Pairwise comparisons of estimated search cost CDFs.



Note: Graphs present estimated search cost distributions in a way that allows to compare them across one characteristic (see top of each column), keeping all the other ones fixed at all their possible values (see graph headings).

References

- ACKERBERG, D., C. L. BENKARD, S. BERRY, AND A. PAKES (2007): “Econometric Tools for Analysing Market Outcomes,” in *Handbook of Econometrics*, ed. by J. J. Heckman and E. Leamer, Elsevier B.V., vol. 6A, chap. 63, 4171–4276.
- AGARWAL, S., J. GRIGSBY, A. HORTAÇSU, G. MATVOS, A. SERU, AND V. YAO (2017): “Search and Screening in Credit Markets Search and Screening in Credit Markets,” Working Paper.
- AGRAWAL, D. (1996): “Effect of Brand Loyalty on Advertising and Trade Promotions: A Game Theoretic Analysis with Empirical Evidence,” *Marketing Science*, 15, 86–108.
- AGUIRREGABIRIA, V. (1999): “The Dynamics of Markups and Inventories in Retailing Firms,” *Review of Economic Studies*, 66, 275–308.
- AGUIRREGABIRIA, V. AND P. MIRA (2007): “Sequential Estimation of Dynamic Discrete Games,” *Econometrica*, 75, 1–53.
- AGUIRREGABIRIA, V. AND A. NEVO (2013): “Recent Developments in Empirical IO: Dynamic Demand and Dynamic Games,” in *Advances in Economics and Econometrics: Tenth World Congress*, ed. by D. Acemoglu, M. Arellano, and E. Dekel, Cambridge University Press, vol. 3: Econometrics, 53–122.
- AGUIRREGABIRIA, V. AND J. SUZUKI (2014): “Identification and Counterfactuals in Dynamic Models of Market Entry and Exit,” *Quantitative Marketing and Economics*, 12, 267–304.
- AILAWADI, K. L., B. A. HARLAM, J. CÉSAR, AND D. TROUNCE (2006): “Promotion Profitability for a Retailer: The Role of Promotion, Brand, Category and Store Characteristics,” *Journal of Marketing Research*, 43, 518–535.
- ALEXANDROV, A. AND S. KOULAYEV (2018): “No Shopping in the U.S. Mortgage Market:

- Direct and Strategic Effects of Providing Information,” CFPB Office of Research Working Paper No. 2017-01.
- ALLEN, J., R. CLARK, AND J.-F. HOUDE (2013): “The Effect of Mergers in Search Markets: Evidence from the Canadian Mortgage Industry,” *American Economic Review*, 104, 3365–3396.
- (2014): “Price Dispersion in Mortgage Markets,” *Journal of Industrial Economics*, 62, 377–416.
- (2017): “Search Frictions and Market Power in Negotiated Price Markets,” *Journal of Political Economy*, forthcoming.
- ANDERSON, S. P. AND A. DE PALMA (2006): “Market Performance with Multiproduct Firms,” *Journal of Industrial Economics*, 54, 95–124.
- ARCIDIACONO, P. AND P. B. ELLICKSON (2011): “Practical Methods for Estimation of Dynamic Discrete Choice Models,” *Annual Review of Economics*, 3, 363–394.
- ARIE, G. AND P. L. E. GRIECO (2014): “Who pays for switching costs?” *Quantitative Marketing and Economics*, 12, 379–419.
- ARMSTRONG, M. (2006): “Competition in Two-Sided Markets,” *RAND Journal of Economics*, 37, 668–691.
- ATHEY, S. AND P. A. HAILE (2007): “Nonparametric Approaches to Auctions,” in *Handbook of Econometrics*, ed. by E. E. L. James J. Heckman, Elsevier, vol. 6A, 3847–3965.
- BAJARI, P., C. L. BENKARD, AND J. LEVIN (2007): “Estimating Dynamic Models of Imperfect Competition,” *Econometrica*, 75, 1331–1370.
- BARWICK, P. J. AND P. A. PATHAK (2015): “The Costs of Free Entry: An Empirical Study of Real Estate Agents in Greater Boston,” *RAND Journal of Economics*, 46, 103–145.
- BASKER, E. (2012): “Raising the Barcode Scanner: Technology and Productivity in the Retail Sector,” *American Economic Journal: Applied Economics*, 4, 1–27.
- (2015): “Change at the Checkout: Tracing the Impact of a Process Innovation,” *The Journal of Industrial Economics*, 63, 339–370.
- BAYE, M. R. AND J. MORGAN (2001): “Information Gatekeepers on the Internet and the Competitiveness of Homogeneous Product Markets,” *American Economic Review*, 91, 454–474.

- BAYE, M. R., J. MORGAN, AND P. SCHOLTEN (2006): “Information, Search, and Price Dispersion,” in *Handbook of Economics and Information Systems*, ed. by T. Hendershott, Elsevier Press, Amsterdam.
- BEGGS, A. AND P. KLEMPERER (1992): “Multi-Period Competition with Switching Costs,” *Econometrica*, 60, 651–66.
- BENABOU, R. (1993): “Search Market Equilibrium, Bilateral Heterogeneity, and Repeat Purchases,” *Journal of Economic Theory*, 60, 140–158.
- BENETTON, M. (2018): “Leverage Regulation and Market Structure: An Empirical Model of the UK Mortgage Market,” Working Paper.
- BERGSTRESSER, D., J. M. CHALMERS, AND P. TUFANO (2007): “Assessing the Costs and Benefits of Brokers in the Mutual Fund Industry,” *Review of Financial Studies*, 22, 4129–4156.
- BERRY, S. AND P. HAILE (2014): “Identification in Differentiated Products Markets,” *Econometrica*, 82, 1749–1797.
- BERRY, S., J. LEVINSOHN, AND A. PAKES (1995): “Automobile Prices in Market Equilibrium,” *Econometrica*, 63, 841–890.
- BIGLAISER, G. (1993): “Middlemen as experts,” *RAND Journal of Economics*, 24, 212–223.
- BLATTBERG, R. C. AND R. A. BRIESCH (2010): “Sales Promotions,” in *Oxford Handbook of Pricing Management*, Oxford University Press.
- BLOOM, P. N., G. T. GUNDLACH, AND J. P. CANNON (2000): “Slotting Allowances and Fees: Schools of Thought and the Views of Practicing Managers,” *Journal of Marketing*, 64, 92–108.
- BURDETT, K. AND K. JUDD (1983): “Equilibrium Price Dispersion,” *Econometrica*, 51, 955–969.
- CHEN, S. X. (1999): “Beta Kernel Estimators for Density Functions,” *Computational Statistics & Data Analysis*, 31, 131–145.
- CHINTAGUNTA, P. (2002): “Investigating Category Pricing Behavior at a Retail Chain,” *Journal of Marketing Research*, 39, 141–154.
- COMPETITION COMMISSION (2000): “Supermarkets: A Report on the Supply of Groceries from Multiple Stories in the United Kingdom,” Tech. Rep. CM 4882.

- CONLISK, J., E. GERSTNER, AND J. SOBEL (1984): “Cyclic Pricing by a Durable Goods Monopolist,” *Quarterly Journal of Economics*, 99, 489–505.
- DE CORNIÈRE, A. AND G. TAYLOR (2017): “A Model of Biased Intermediation,” Working Paper, Toulouse School of Economics.
- DE LOS SANTOS, B., A. HORTAÇSU, AND M. R. WILDENBEEST (2012): “Testing Models of Consumer Search Using Data on Web Browsing and Purchasing behavior,” *American Economic Review*, 102, 2455–2480.
- DELTAS, G. AND Z. LI (2018): “Free Riding on the Search of Others: Information Externalities in the Mortgage Industry,” Working Paper, University of Illinois.
- DEUFLHARD, F. (2016): “Quantifying Inertia in Retail Deposit Markets,” Working Paper.
- DORASZELSKI, U. AND M. SATTERTHWAITE (2010): “Computable Markov-Perfect Industry Dynamics,” *RAND Journal of Economics*, 41, 215–243.
- DRISKILL, R. A. AND S. MCCAFFERTY (1989): “Dynamic Duopoly with Adjustment Costs: A Differential Game Approach,” *Journal of Economic Theory*, 49, 324–338.
- DUBÉ, J.-P., G. J. HITSCH, AND P. E. ROSSI (2009): “Do Switching Costs Make Markets Less Competitive?” *Journal of Marketing Research*, 46, 435–445.
- (2010): “State dependence and alternative explanations for consumer inertia,” *RAND Journal of Economics*, 41, 417–445.
- DUBÉ, J.-P., G. J. HITSCH, P. E. ROSSI, AND M. A. VITORINO (2008): “Category Pricing with State-Dependent Utility,” *Marketing Science*, 27, 417–429.
- DUBOIS, P., R. GRIFFITH, AND A. NEVO (2014): “Do Prices and Attributes Explain International Differences in Food Purchases?” *American Economic Review*, 104, 832–867.
- DUTTA, S., M. BERGEN, D. LEVY, AND R. VENABLE (1999): “Menu Costs, Posted Prices, and Multiproduct Retailers,” *Journal of Money, Credit and Banking*, 31, 683–703.
- EDELMAN, B. AND J. WRIGHT (2015): “Price Coherence and Excessive Intermediation,” *Quarterly Journal of Economics*, 130, 1283–1328.
- EGAN, M. (2018): “Brokers vs. Retail Investors: Conflicting Interests and Dominated Products,” *Journal of Finance*, forthcoming.

- EGAN, M., G. MATVOS, AND A. SERU (2018): “The Market for Financial Adviser Misconduct,” *Journal of Political Economy*, forthcoming.
- EIZENBERG, A. AND A. SALVO (2015): “The Rise of Fringe Competitors in the Wake of an Emerging Middle Class: An Empirical Analysis,” *American Economic Journal: Applied Economics*, 7, 85–122.
- ELLISON, S. F., C. M. SNYDER, AND H. ZHANG (2015): “Costs of Managerial Attention and Activity as a Source of Sticky Prices: Structural Estimates from an Online Market,” Working Paper, MIT.
- ERICSON, R. AND A. PAKES (1995): “Markov-Perfect Industry Dynamics: A Framework for Empirical Work,” *Review of Economic Studies*, 62, 53–82.
- FARRELL, J. AND P. KLEMPERER (2007): “Coordination and Lock-In: Competition with Switching Costs and Network Effects,” in *Handbook of Industrial Organization*, ed. by M. Armstrong and R. Porter, Elsevier B.V., vol. 3, chap. 31, 1967–2072.
- FERSHTMAN, C. AND M. I. KAMIEN (1987): “Dynamic Duopolistic Competition with Sticky Prices,” *Econometrica*, 55, 1151–1164.
- FERSHTMAN, C. AND A. PAKES (2012): “Dynamic Games with Asymmetric Information: A Framework for Empirical Work,” *Quarterly Journal of Economics*, 127, 1611–61.
- FOWLIE, M., M. REGUANT, AND S. P. RYAN (2016): “Market-Based Emissions Regulation and Industry Dynamics,” *Journal of Political Economy*, 124, 249–302.
- FRANKEL, A. S. (1998): “Monopoly and Competition in the Supply and Exchange of Money,” *Antitrust Law Journal*, 66, 313–361.
- GALEOTTI, A. AND J. L. MORAGA-GONZÁLEZ (2009): “Platform Intermediation in a Market for Differentiated Products,” *European Economic Review*, 53, 417–428.
- GAVAZZA, A. (2016): “An Empirical Equilibrium Model of a Decentralized Asset Market,” *Econometrica*, 84, 1755–1798.
- GENTRY, M., T. LI, AND J. LU (2015): “Identification and Estimation in First-Price Auctions with Risk-Averse Bidders and Selective Entry,” Working Paper, LSE.
- GOETTLER, R. L. AND B. R. GORDON (2011): “Does AMD Spur Intel to Innovate More?” *Journal of Political Economy*, 119, 1141–1200.

- GRIFFITH, R., E. LEIBTAG, A. LEICESTER, AND A. NEVO (2009): “Consumer Shopping Behavior: How Much Do Consumers Save?” *Journal of Economic Perspectives*, 23, 99–120.
- GRIFFITH, R., L. NESHEIM, AND M. O’CONNELL (2017): “Income effects and the welfare consequences of tax in differentiated product oligopoly,” *Quantitative Economics*, forthcoming.
- GUERRE, E., I. PERRIGNE, AND Q. VUONG (2000): “Optimal Nonparametric Estimation of First-Price Auctions,” *Econometrica*, 68, 525–574.
- GUIMARAES, B. AND K. D. SHEEDY (2011): “Sales and Monetary Policy,” *American Economic Review*, 101, 844–876.
- GUIO, L., A. POZZI, A. TSOY, L. GAMBACORTA, AND P. MISTRULLI (2017): “The Cost of Distorted Financial Advice: Evidence from the Mortgage Market,” EIEF Working Paper 17/13.
- HAILE, P. A. AND Y. KITAMURA (2018): “Unobserved Heterogeneity in Auctions,” Cowles Foundation Discussion Paper No. 2141.
- HALL, G. AND J. RUST (2003): “Middle Men versus Market Makers: A Theory of Competitive Exchange,” *Journal of Political Economy*, 111, 353–403.
- HANSEN, B. (2008): “Uniform Convergence Rates for Kernel Estimation with Dependent Data,” *Econometric Theory*, 24, 726–748.
- HÄRDLE, W. (1991): *Smoothing Techniques with Implementation in S*, New York: Springer Verlag.
- HASTINGS, J. S., B. C. MADRIAN, AND W. L. SKIMMYHORN (2013): “Financial Literacy, Financial Education and Economic Outcomes,” *Annual Review of Economics*, 1, 347–373.
- HAYFIELD, T. AND J. S. RACINE (2008): “Nonparametric Econometrics: The np Package,” *Journal of Statistical Software*, 27.
- HEIDHUES, P. AND B. KŐSZEGI (2014): “Regular prices and sales,” *Theoretical Economics*, 9, 217–251.
- HENDEL, I., A. LIZZERI, AND N. ROKETSKIY (2014): “Nonlinear Pricing of Storable Goods,” *American Economic Journal: Microeconomics*, 6, 1–34.
- HENDEL, I. AND A. NEVO (2013): “Intertemporal Price Discrimination in Storable Goods Markets,” *American Economic Review*, 103, 2722–2751.

- HICKMAN, B. R. AND T. P. HUBBARD (2015): “Replacing Sample Trimming with Boundary Correction in Nonparametric Estimation of First-Price Auctions,” *Journal of Applied Econometrics*, 30, 739–762.
- HONG, H. AND M. SHUM (2006): “Using Price Distribution to Estimate Search Costs,” *RAND Journal of Economics*, 37, 257–275.
- HONKA, E. AND P. CHINTAGUNTA (2017): “Simultaneous or Sequential? Search Strategies in the US Auto Insurance Industry,” *Marketing Science*, 36, 21–42.
- HONKA, E., A. HORTAÇSU, AND M. A. VITORINO (2017): “Advertising, Consumer Awareness and Choice: Evidence from the U.S. Banking Industry,” *RAND Journal of Economics*, 48, 611–646.
- HORSKY, D., P. PAVLIDIS, AND M. SONG (2012): “Incorporating State Dependence in Aggregate Brand-level Demand Models,” Working Paper, Simon Graduate School of Business, University of Rochester.
- HOSKEN, D. AND D. REIFFEN (2004): “Patterns of Retail Price Variation,” *RAND Journal of Economics*, 35, 128–146.
- HOTZ, V. J. AND R. A. MILLER (1993): “Conditional Choice Probabilities and the Estimation of Dynamic Models,” *Review of Economic Studies*, 60, 497–529.
- HOTZ, V. J., R. A. MILLER, S. SANDERS, AND J. SMITH (1994): “A Simulation Estimator for Dynamic Models of Discrete Choice,” *Review of Economic Studies*, 61, 265–289.
- ICHIMURA, H. (1993): “Semiparametric Least Squares (SLS) and Weighted SLS Estimation of Single-Index Models,” *Journal of Econometrics*, 58, 71–120.
- IMLA (2015): “The Changing Face of Mortgage Distribution,” Intermediary Mortgage Lenders Association report.
- INDERST, R. AND M. OTTAVIANI (2012a): “Competition through Commissions and Kickbacks,” *American Economic Review*, 102, 780–809.
- (2012b): “Financial Advice,” *Journal of Economic Literature*, 50, 494–512.
- (2012c): “How (Not) to Pay for Advice: A Framework for Consumer Financial Protection,” *Journal of Financial Economics*, 105, 393–411.

- ISCENKO, Z. (2018): “Choices of Dominated Mortgage Products by UK Consumers,” FCA Occasional Paper No. 33.
- ISCENKO, Z. AND J. NIEBOER (2018): “Effects of the advice requirement and intermediation in the UK mortgage market,” FCA Occasional Paper No. 34.
- JOFRE-BONET, M. AND M. PESENDORFER (2003): “Estimation of a Dynamic Auction Game,” *Econometrica*, 71, 1443–1489.
- JOHNSON, J. P. AND D. P. MYATT (2006): “Multiproduct Cournot oligopoly,” *RAND Journal of Economics*, 37, 583–601.
- JUN, B. AND X. VIVES (2003): “Strategic incentives in dynamic duopoly,” *Journal of Economic Theory*, 116, 249–281.
- KADIYALI, V., P. CHINTAGUNTA, AND N. VILCASSIM (2000): “Manufacturer-Retailer Channel Interactions and Implications for Channel Power: An Empirical Investigation of Pricing in a Local Market,” *Marketing Science*, 19, 127–248.
- KANO, K. (2013): “Menu Costs and Dynamic Duopoly,” *International Journal of Industrial Organization*, 31, 102–118.
- KARUNAMUNI, R. J. AND S. ZHANG (2008): “Some Improvements on a Boundary Corrected Kernel Density Estimator,” *Statistics & Probability Letters*, 78, 499–507.
- KASHYAP, A. K. AND M. ROSTOM (2018): “Poor Mortgage Choice,” Working Paper.
- KIM, M., D. KLIGER, AND B. VALE (2003): “Estimating Switching Costs: The Case of Banking,” *Journal of Financial Intermediation*, 12, 25–56.
- KOMAROVA, T., F. A. M. SANCHES, D. SILVA JR., AND S. SRISUMA (forthcoming): “Joint Analysis of the Discount Factor and Payoff Parameters in Dynamic Discrete Choice Models,” *Quantitative Economics*.
- LAL, R. (1990): “Price Promotions: Limiting Competitive Encroachment,” *Marketing Science*, 9, 247–262.
- LAPHAM, B. AND R. WARE (1994): “Markov puppy dogs and related animals,” *International Journal of Industrial Organization*, 12, 569–593.
- LEE, C.-Y., J. W. ROBERTS, AND A. SWEETING (2012): “Competition and Dynamic Pricing in a Perishable Goods Market,” Working Paper, Duke University.

- LEVY, D., M. BERGEN, S. DUTTA, AND R. VENABLE (1997): “The Magnitude of Menu Costs: Direct Evidence from Large U.S. Supermarket Chains,” *Quarterly Journal of Economics*, 112, 791–825.
- LI, H. AND N. LIU (2015): “Nonparametric Identification of Double Auctions with Bargaining,” Working Paper, Shanghai University of Finance and Economics.
- MACMINN, R. D. (1980): “Search and Market Equilibrium,” *Journal of Political Economy*, 88, 308–327.
- MARRON, J. S. AND D. RUPPERT (1994): “Transformations to Reduce Boundary Bias in Kernel Density Estimation,” *Journal of the Royal Statistical Society Series B*, 56, 653–671.
- MASRY, E. (1996): “Multivariate Local Polynomial Regression for Time Series: Uniform Strong Consistency and Rates,” *Journal of Time Series Analysis*, 17, 571–599.
- MATĚJKA, F. (2016): “Rationally Inattentive Seller: Sales and Discrete Pricing,” *Review of Economic Studies*, forthcoming, cERGE-EI Working Paper.
- MCAFEE, R. P. (1995): “Multiproduct Equilibrium Price Dispersion,” *Journal of Economic Theory*, 67, 83–105.
- MEZA, S. AND K. SUDHIR (2010): “Do private labels increase retailer bargaining power?” *Quantitative Marketing and Economics*, 8, 333–363.
- MILGROM, P. AND I. SEGAL (2002): “Envelope Theorems for Arbitrary Choice Sets,” *Econometrica*, 70, 583–601.
- MILGROM, P. R. AND R. J. WEBER (1985): “Distributional Strategies for Games with Incomplete Information,” *Mathematics of Operations Research*, 10, 619–632.
- MORAGA-GONZÁLEZ, J. L., Z. SÁNDOR, AND M. R. WILDENBEEST (2013): “Semi-Nonparametric Estimation of Consumer Search Costs,” *Journal of Applied Econometrics*, 28, 1205–1223.
- (2017): “Non-sequential Search Equilibrium with Search Cost Heterogeneity,” *International Journal of Industrial Organization*, 50, 392–414.
- MORAGA-GONZÁLEZ, J. L. AND M. R. WILDENBEEST (2007): “Maximum Likelihood Estimation of Search Costs,” *European Economic Review*, 52, 820–848.

- NAKAMURA, E. AND J. STEINSSON (2008): “Five Facts About Prices: A Reevaluation Of Menu Cost Models,” *Quarterly Journal of Economics*, 123, 1415–1464.
- NAVA, F. AND P. SCHIRALDI (2014): “Sales and Collusion in a Market with Storage,” *Journal of the European Economic Association*, 12, 791–832.
- NEVO, A. (2001a): “Measuring Market Power in the Ready-To-Eat Cereal Industry,” *Econometrica*, 69, 307–342.
- (2001b): “Mergers with Differentiated Products: the Case of the Ready-to-Eat Cereal Industry,” *RAND Journal of Economics*, 31, 395–421.
- PAVLIDIS, P. AND P. ELLICKSON (2017): “Implications of Parent Brand Inertia for Multiproduct Pricing,” *Quantitative Marketing and Economics*, 15, 188–227.
- PEARCY, J. (2014): “Bargains Followed by Bargains: When Switching Costs Make Markets More Competitive,” Working Paper, Montana State University.
- PEREIRA, P. (2005): “Multiplicity of Equilibria In Search Markets with Free Entry and Exit,” *International Journal of Industrial Organization*, 23, 325–339.
- PESENDORFER, M. (2002): “Retail Sales: A Study of Pricing Behavior in Supermarkets,” *Journal of Business*, 75, 33–66.
- PESENDORFER, M. AND P. SCHMIDT-DENGLER (2008): “Asymptotic Least Squares Estimators for Dynamic Games,” *Review of Economic Studies*, 75, 901–928.
- POWELL, J., J. STOCK, AND T. STOKER (1989): “Semiparametric Estimation of Index Coefficients,” *Econometrica*, 57, 1403–1430.
- RENY, P. J. (2011): “On the Existence of Monotone Pure-Strategy Equilibria in Bayesian games,” *Econometrica*, 79, 499–553.
- RHODES, A. (2014): “Multiproduct Retailing,” *Review of Economic Studies*, forthcoming.
- ROCHET, J.-C. AND J. TIROLE (2006): “Two-Sided Markets: A Progress Report,” *RAND Journal of Economics*, 37, 645–667.
- ROSENTHAL, R. W. (1980): “A Model in Which an Increase in the Number of Sellers Leads to a Higher Price,” *Econometrica*, 48, 1575–1580.

- RUBINSTEIN, A. AND A. WOLINSKY (1987): "Middlemen," *Quarterly Journal of Economics*, 102, 581–594.
- SALOP, S. C. AND J. E. STIGLITZ (1977): "Bargains and Ripoffs: A Model of Monopolistically Competitive Price Dispersion," *Review of Economic Studies*, 44, 493–510.
- SALZ, T. (2017): "Intermediation and Competition in Search Markets: An Empirical Case Study," Working Paper, Columbia University.
- SANCHES, F., D. SILVA JUNIOR, AND S. SRISUMA (2016a): "Minimum Distance Estimation of Search Costs using Price Distribution," *Journal of Business and Economic Statistics*, forthcoming.
- SANCHES, F. A. M., D. SILVA JR., AND S. SRISUMA (2016b): "Ordinary Least Squares Estimation of a Dynamic Game Model," *International Economic Review*, forthcoming.
- SCHMALENSEE, R. (1976): "A Model of Promotional Competition in Oligopoly," *Review of Economic Studies*, 43, 493–507.
- SEILER, S. (2013): "The impact of search costs on consumer behavior: A dynamic approach," *Quantitative Marketing and Economics*, 11, 155–203.
- SHELEGIA, S. (2012): "Multiproduct Pricing in Oligopoly," *International Journal of Industrial Organization*, 30, 231–242.
- SLADE, M. E. (1998): "Optimal Pricing with Costly Adjustment: Evidence from Retail Grocery Prices," *Review of Economic Studies*, 65, 87–107.
- (1999): "Sticky prices in a dynamic oligopoly: An investigation of (s, S) thresholds," *International Journal of Industrial Organization*, 17, 477–511.
- SOBEL, J. (1984): "The Timing of Sales," *Review of Economic Studies*, 51, 353–368.
- SPULBER, D. F. (1995): "Bertrand Competition when Rivals' Costs are Unknown," *Journal of Industrial Economics*, 43, 1–11.
- SRINIVASAN, S., K. PAUWELS, D. M. HANSENS, AND M. G. DEKIMPE (2004): "Do Promotions Benefit Manufacturers, Retailers, or Both?" *Management Science*, 50, 617–629.
- STONE, C. J. (1982): "Optimal Rate of Convergence for Nonparametric Regressions," *Annals of Statistics*, 10, 1040–1053.

- SWEETING, A. (2012): "Price Dynamics in Perishable Goods Markets: The Case of Secondary Markets for Major League Baseball Tickets," *Journal of Political Economy*, 120, 1133–1172.
- (2013): "Dynamic Product Positioning in Differentiated Product Markets: The Effect of Fees for Musical Performance Rights on the Commercial Radio Industry," *Econometrica*, 81, 1763–1803.
- VARIAN, H. R. (1980): "A Model of Sales," *American Economic Review*, 70, 651–659.
- VILLAS-BOAS, J. M. (2015): "A Short Survey on Switching Costs and Dynamic Competition," *International Journal of Research in Marketing*, forthcoming.
- VIVES, X. (2002): *Oligopoly pricing: old ideas and new tools*, The MIT Press, Cambridge, MA.
- WAND, M. P. AND M. C. JONES (1999): *Kernel Smoothing*, Chapman & Hall/CRC Monographs on Statistics & Applied Probability.
- WOODWARD, S. E. AND R. E. HALL (2012): "Diagnosing Consumer Confusion and Sub-optimal Shopping Effort: Theory and Mortgage Market Evidence," *American Economic Review*, 102, 3249–3276.
- YAVAŞ, A. (1994): "Middlemen in Bilateral Search Markets," *Journal of Labor Economics*, 12, 406–429.
- ZBARACKI, M. J., M. RITSON, D. LEVY, S. DUTTA, AND M. BERGEN (2004): "Managerial and Customer Costs of Price Adjustment: Direct Evidence from Industrial Markets," *Review of Economics and Statistics*, 86, 514–533.