

Essays in Industrial Organisation

Price Competition, Strategic Obfuscation, Advertising & Consumer Behaviour



UNIVERSITY OF
LIVERPOOL

Thesis submitted in accordance with the requirements of
The University of Liverpool for the degree of
Doctor of Philosophy

by

Robert Andrew Edwards

Department of Economics, Finance & Accounting

January 2017

This thesis is dedicated to
Rebekah McCarthy,
Frank & Mary Young.

Acknowledgements

There are several individuals to whom I must extend my sincere gratitude. Firstly, to Professor Gary Cook, Head of Department of Economics, Finance and Accounting, for inspiring me to pursue postgraduate studies. Gary has supported my development throughout my academic education. His enthusiasm, encouragement and direction will have a lasting impact on my approach to both research and teaching.

Secondly, I am indebted to my advisors: Dr. Yiquan Gu, Professor Brendan McCabe and Dr. Yao Rao. Yiquan's expertise in the application of Game Theory to topics in Industrial Organisation stimulated my interest in this research area. Brendan's wealth of knowledge and experience has been invaluable in developing this thesis.

Thirdly, I must thank the members of the Microeconomics Research Group at the University of Liverpool Management School. Special thanks are due to Professor Dominique Demougin and to Dr. Robert Routledge. I must also thank my friends and doctoral colleagues: Claire, Sam, Michael, Yazan, Christiane, Harvey & Sanna.

The research articles contained in this thesis have also benefited from comments at several international conferences including: North West Doctoral Training Centre Annual Conference (2015), Bielefeld Doctoral School Workshop on Economic Theory (2015), Workshop on Game Theory & Social Choice at the University of Hungary (2015) and the Annual Royal Economic Society Conference held at the University of Brighton (2016). I also gratefully acknowledge the financial support of the Economic & Social Research Council.

Finally, I must thank Rebekah and my family for their constant support throughout my academic studies, without whom the production of this thesis would not have been possible.

Contents

1	Introduction	1
2	Obfuscation, Strategic Pricing & Consumer Behaviour.	4
2.1	Introduction	4
2.2	Theory	6
2.2.1	Multi-Dimensional Pricing & Composite Products	6
2.2.1.1	Endogenous Consumer Sophistication	10
2.2.2	Price Framing Effects	11
2.2.3	Saliency	18
2.2.3.1	Attention Grabbing Products	23
2.2.4	Search Costs	25
2.2.4.1	Multi-Product Retailing	28
2.2.5	Frequently Changing Prices	35
2.3	Evidence	37
2.3.1	Empirical Evidence	37
2.3.2	Experiments	45
2.3.2.1	Differentiated Products	51
2.3.2.2	Natural Increases in Choice	57
2.3.2.3	Saliency	59
2.4	Simplicity Biases: Theory & Evidence	61
2.5	Conclusion	73
3	Competition with Simplicity Biased Consumers	75
3.1	Introduction	75
3.1.1	Summary of Main Results	78
3.1.2	Related Literature	79
3.2	A Model of Competition with Simplicity Biased Consumers	82
3.3	Equilibrium Pricing	84
3.3.1	Symmetric Pricing	85

3.3.2	Asymmetric Pricing	87
3.4	Equilibrium Obfuscation	90
3.4.1	Symmetric Equilibria	90
3.4.2	Asymmetric Equilibria	92
3.4.3	Mixed Strategy Equilibrium	93
3.5	Welfare Analysis	95
3.6	Discussion	96
3.6.1	The Role of Timing	100
3.6.2	Divergent Consumer Preferences	100
3.6.3	Costly Obfuscation	100
3.6.4	Further Extensions	101
3.7	Conclusion	102
3.8	Appendix A: Main Proofs	103
3.9	Appendix B: Alternative Characterisation of Firm Profit.	110
3.10	Appendix C: Supplementary Material	111
4	Promotions with Biased Consumers	113
4.1	Introduction	113
4.1.1	Related Literature	117
4.2	The Model	120
4.2.1	Equilibrium Payoffs	123
4.2.2	Equilibrium Obfuscation	126
4.2.3	Derestricting the Simplicity Bias.	127
4.2.4	Firm Profit & Price Distributions	127
4.2.5	Decreasing Marginal Effect of Obfuscation	128
4.3	Comparative Statics & Policy Measures	129
4.3.1	Comparative Statics	129
4.3.2	Policies to Improve Consumer Sophistication	130
4.3.3	An Illustrative Example	132
4.4	Endogenous Simplicity Biases	134
4.4.1	Strategic Advertising	134
4.4.2	False Advertising	136
4.5	Discussion	137
4.5.1	Policy Implications	140
4.5.2	Hypotheses for Experimental & Empirical Research	140
4.5.3	Limitations & Extensions	141

4.6	Conclusion	142
4.7	Appendix A: Main Proofs	143
4.8	Appendix B: Supplementary Materials.	151
5	Conclusion & Further Research	157
5.1	Conclusions	157
5.2	Further Research Projects	161
5.2.1	Low Income Consumers and Add-On Insurance Products	161
5.2.2	Commercial Obfuscation	162
	Bibliography	163

List of Figures

2.1	Hossain & Morgan (2006)	45
2.2	Treatment Types: Taiwan, Brown et al (2010).	47
2.3	Treatment Types: Ireland, Brown et al (2010).	47
2.4	Kalayci and Potters (2011).	52
3.1	Symmetric Pricing: $\{L, L\}$	89
3.2	Symmetric Pricing: $\{H, H\}$	89
3.3	Asymmetric Pricing: $\{L, H\}$	89
3.4	Asymmetric Pricing: $\{L, H\}$	89
3.5	Equilibrium Obfuscation I	94
3.6	Equilibrium Obfuscation II	94
4.1	Equilibrium Pricing: $\delta < \delta^c$	123
4.2	Probability Assigned to Low Complexity	132
4.3	Price Distribution ($\delta = 0.25$)	133
4.4	Price Distribution ($\delta = 0.4$)	133

Chapter 1

Introduction

This thesis explores the incentives for firms to introduce unnecessary complexity into their pricing schemes or product information with the intention of confusing consumers, a practice that industrial economists refer to as strategic obfuscation. This type of behaviour includes firms decomposing their price into excessively many components, using non-standard terminology or measurements to express their price, using mathematically complex prices and updating their prices frequently. Obfuscation also includes the ability of firms to increase the difficulty, time or effort required for consumers to locate a product or its price, within a store or on a website.

Obfuscation is a growing concern for competition authorities. In the UK, the Competition & Markets Authority reports that complex pricing harms retail market competition by suppressing consumer engagement and limiting switching between sellers.¹ Therefore, understanding the precise incentives for a firm to obfuscate, enables competition authorities to formulate policies that safeguard consumers and enhance the competitive process.

This thesis is organised as three related articles with distinct contributions. In Chapter 2, I present a comprehensive review of both the theory and evidence relating to obfuscation and price competition. The primary contribution is to provide an overview of this substantial body of literature. The secondary contribution is to highlight several inconsistencies across this research area. This motivates further study and the subsequent chapters contained in this thesis.

In Chapter 2, I also review how consumers tend to respond to obfuscation strategies of firms in retail markets, which remains an under-researched topic. I document strong evidence for a simplicity bias by consumers, according to which buyers actively discriminate against firms that obfuscate their prices. This idea has not been fully explored in the context of strategic obfuscation and draws insights from the related disciplines of Law, Marketing and Psychology.

In Chapter 3, I present a model in which each firm selects a pricing scheme that consists of a price for their product and a complexity level with which to present their price. In the game,

¹www.gov.uk/government/news/cma-proposes-better-deal-for-bank-customers. www.gov.uk/government/news/cma-sets-out-case-for-energy-market-reform. Accessed 03/16.

firms commit to the complexity of their pricing scheme before selecting their price. This captures the idea that prices can be adjusted quickly but the complexity of the pricing scheme requires a longer time to adjust. Therefore, obfuscation in this model might specifically include: The ease with which consumers can navigate a firm's website to identify the price, or the complexity of the terms and conditions attached to the product's price.

One contribution of Chapter 3 is to formally introduce the idea that consumers are biased towards simple pricing schemes and to investigate this idea further than has been considered by previous studies. The central tension borne out in the model is that a firm can soften price competition in the market by obfuscating its prices, which prevents more consumers from identifying the seller with the lowest price. This incentivises firms to obfuscate their prices. However, confused consumers prefer to purchase from the firm with the least complex pricing scheme. Therefore, the seller that obfuscates most will obtain a smaller fraction of the consumers who are confused. This places downward pressure on the incentive for each firm to obfuscate.

In equilibrium, this incentive structure induces a mixed strategy in both prices and obfuscation for firms. Therefore, sellers randomise both their price and the complexity of their price according to well-defined probability distribution functions. This is consistent with the behaviour of firms in modern retail markets. In addition, the firm with the least complex pricing structure will charge the highest price, to capitalise on consumers' preferences for simple pricing. If the simplicity biases of confused consumers are removed from the model, firms universally select maximum obfuscation and follow a mixed strategy in pricing their product.

The impact of simplicity biased consumers on consumer welfare is particularly novel and interesting. Simplicity biased consumers buy from the firm with the lowest obfuscation and this firm charges the highest prices, which increases firm profit. However, the presence of simplicity biased consumers exerts downward pressure on the incentives for firms to obfuscate. This leads to lower obfuscation and enables more consumers to understand prices. Price competition subsequently intensifies and firm profit decreases. It is demonstrated that the presence of simplicity biased consumers, who are biased towards the simplest and most expensive firm, can reduce average prices in the market and increase consumer welfare. In this direction, Chapter 3 investigates the implications of policies to inform consumers that firms may attach a price premium to simple pricing schemes. I show that such policies can often leave consumers worse off.

In Chapter 4, I consider a restatement of the model developed in Chapter 3. The primary difference is that firms select their price and obfuscation level simultaneously. This captures market settings in which obfuscation and prices are both instantaneously adjustable. Therefore, obfuscation in this framework includes: The quantity of dimensions used in a pricing scheme and the complexity of the units used to express the total price.

In the fully simultaneous framework in Chapter 4, the firm with the highest price chooses the most complex price scheme. The intuition follows that the least competitive firm benefits most from confusing consumers because consumers who are able to evaluate prices will always buy from the cheaper rival. Therefore, consumers who are biased towards the simplest pricing scheme will obtain a lower price than if they had selected a seller at random. The welfare impact of the simplicity bias is also now unidirectional. Biased consumers purchase from the cheapest seller and their presence continues to suppress firms' obfuscation incentives, which always increases consumer welfare.

In a series of extensions, I develop the model to explore the implications of several consumer protection and competition policies. This includes policies that influence consumers' aversion to complex prices, and policies that enable more consumers to compare the prices available for any given level of market complexity. The strength of consumers' preference for simple pricing, termed the degree of simplicity bias, is also endogenised as a function of both regulatory policy and firms' advertising campaigns. This sharpens our understanding of the circumstances in which it is profitable for firms to stimulate and suppress consumers' biases.

In Chapter 4, I also distinguish between obfuscation and advertising. This distinction is not present in mainstream models of obfuscation. The strategy for each firm now consists of three dimensions: Pricing, obfuscation and advertising. Obfuscation determines the fraction of consumers who are able to compare prices. Advertising determines whether consumers perceive a firm's price to be simple or complex. Therefore, a firm in the extended model is able to combine an intrinsically complicated price scheme that is difficult for consumers to compare, with advertising that makes the price scheme appear simple. I show that this distinction elicits a novel form of false advertising that has become a common practice in retail markets but is absent from academic research.

In Chapter 5, I conclude this thesis by discussing the overarching contributions of the models in the context of the wider literature. I also outline several specific avenues for further research that builds directly on the work reported in this thesis.

Chapter 2

Obfuscation, Strategic Pricing & Consumer Behaviour.

2.1 Introduction

In this chapter I survey the literature on strategic obfuscation, price competition and consumer behaviour in retail markets, including both theory and evidence. The primary objective is to provide an overview of the central themes and connections between these topics of research within Industrial Economics. The secondary objective is to highlight elements of conflict, either in terms of the predictions of the models or the assumptions underpinning the frameworks. To achieve this, I focus in detail on a wide variety of the most influential studies, rather than incorporating every relevant article in this growing field of research. This article is organised with three distinct sections: Theory (Section 2.2), Evidence (Section 2.3) and Simplicity Biases (Section 2.4).

The theoretical literature discussed in Section 2.2 is stratified into five main streams, depending on which ‘imperfections’ are present in the consumer’s decision-making process. Firstly, consumers may neglect some dimensions of a product, which incentivises firms to compete only on the attributes that consumers pay attention to (Section 2.2.1). Secondly, consumers may be limited in their ability to evaluate prices when they are presented in a complicated or non-standard format. This incentivises firms to individuate their promotions or add complexity to their pricing scheme to confuse consumers and soften price competition (Section 2.2.2).

Thirdly, consumers may weight the attributes of a product differently depending on the environment in which it is presented. This incentivises firms to manipulate the wider retail environment by introducing additional products that distort consumers’ purchasing decisions and direct consumers towards high profit items (Section 2.2.3). Fourthly, consumers may rationally choose not to search all of the the prices available in the market due to the time or effort costs involved. This incentivises firms to inflate the cost of searching for their products to reduce the amount of prices that consumers examine. Alternatively, in a multi-product environment, a firm can manipulate the cost of search for each product in its range, to influence the order in which consumers

search (Section 2.2.4). Fifthly, consumers may have a limited memory and be less able to keep pace with prices that change frequently. This incentivises firms to update their prices to prevent consumers from acquiring market knowledge (Section 2.2.5).

The evidence relating to the joint use of obfuscation and strategic pricing strategies is reviewed in Section 2.3. This strand of research is markedly less developed and has, in some cases, failed to maintain pace with the recent surge in theoretical interest in this research area. The empirical literature is synthesised in Section 2.3.1 and includes real world evidence from several markets including the markets for computing memory modules, driving lessons, domestic retail energy and telecommunication carriers. The experimental evidence is considered in Section 2.3.2 and encompasses both field and laboratory experiments.

In some of the laboratory studies, participants are asked to select one of the alternatives available. In this case, participants are incentivised to select the alternative that they perceive to be optimal using ex post reward schemes. In other experiments, participants are asked to spend a nominal period of time in the laboratory and they are able to buy products that they can use during that period. In this case, the individuals will directly consume the products that they select. The series of field experiments also span a diverse set of markets including online auction platforms and retirement financial products. One purpose of reviewing the evidence literature is to highlight areas of accord and departure between real world observations, and the predictions and assumptions of theoretical models.

In Section 2.4, I complete this review by considering the validity of one of the most frequently applied assumptions within this literature: Consumers who cannot compare prices will select a seller at random. This is especially relevant for a body of literature that allows firms to choose the complexity of their prices and promotions. I emphasise that this random purchase rule does not fit particularly well with modern retail markets. This motivates more sophisticated specifications of consumer behaviour, which is a central theme of this thesis and the motivation for the theoretical models contained in Chapter's 3 and 4.

Throughout this article, I focus on non-cooperative games in which firms determine their strategies independently. Special emphasis is placed on two aspects of the underlying model. Firstly, I state precisely the timing of each game because adjustments in timing can induce significant changes for the resulting equilibria. Secondly, I explore in detail the assumptions underpinning consumer behaviour. In subsequent sections of this review and throughout this thesis, I show that some of the most frequently applied assumptions are not well grounded in real world consumer behaviour.

Related reviews include the examination of bounded rationality in Industrial Organisation by Ellison (2006) and the textbook treatment provided by Spiegler (2011). Fudenberg (2006) reviews several relevant advances in behavioural economics and Köszegi (2014) reviews the related

progression of behavioural contract theory. Grubb (2015a) reviews the theoretical and empirical evidence for consumers to make sub-optimal decisions and Grubb (2015b) surveys the implications of consumer overconfidence. Tirole (1988) provides the benchmark textbook for Industrial Organisation theory and Vives (2001) provides a textbook treatment of oligopoly pricing games.

2.2 Theory

2.2.1 Multi-Dimensional Pricing & Composite Products

In many markets, firms sell multiple products that are part of a single purchasing decision for a consumer. For example, the type of printer ink a consumer will buy depends on which printer they have previously purchased (Gabaix and Laibson, 2006). The choice of whether to buy popcorn in a cinema is also conditional on first choosing to visit that particular cinema. Industrial economists refer to these products as composite purchases, where unconstrained consumers should rationally consider every dimension of the product before purchasing any of the items. However, in practice, consumers face many constraints, which leads buyers to neglect some of the product's attributes. This might be the the cost of ink when choosing a printer or the price of popcorn when selecting a cinema.

Consumers may also focus on the price of one product from each seller when comparing firms, before committing to one firm and choosing exactly which product to buy. For example, consumers may compare the basic room rate at hotels but once they have chosen a hotel, they upgrade to a higher quality room without comparing the prices of upgraded rooms elsewhere (Gabaix and Laibson, 2006). In these situations firms have an opportunity to exploit the way that consumers appropriate information.

The seminal model of add-on pricing, due to Ellison (2005), models the idea that firms sell both a basic product and a higher quality upgrade, which consumers may or may not decide to buy depending on their income. Ellison (2005) contrasts a standard pricing game in which sellers advertise all prices to consumers, with an add-on game in which only the price of the base product is advertised. In the add-on game, consumers must visit the store to learn the prices of the add-on (upgrade), which incurs a search cost. There are two types of consumers: High and Low. High types are higher income individuals with a lower marginal utility of income that are assumed to be: (i) Less sensitive to price differences across sellers, (ii) More likely to purchase the higher quality add-on product.

In the standard game, firms simultaneously choose the price of their two products and advertise them to buyers, who subsequently decide which seller to visit. Attending a store incurs a search cost, $s > 0$, for every store visited. Upon visiting, each consumer chooses whether to pur-

chase anything or visit the rival, which incurs a second unit of the search cost. Finally, consumers decide whether to purchase anything from the second store or leave the market empty-handed.

In the add-on game, upgrade prices are hidden. At the beginning, firms choose and advertise only the price of the low quality product. In the second period, firms choose the price of the add-on, conditional on the prices of the low quality good in the market. In the third period, consumers decide which store to visit based solely on the base good price. Upon arriving the consumer learns the price of the higher quality good at that store and decides whether to purchase either product. Consumers may also defer purchasing and visit the rival store, to learn the price of the high quality good.

Ellison (2005) shows that the add-on pricing game creates an adverse selection effect, which limits price competition for low quality products and yields higher profits. The intuition is as follows. Suppose that firms use loss leader pricing strategies where the price of the low quality product is cross subsidised by mark-ups on the higher quality product. A firm may attempt to increase demand for its products by reducing the price of the low quality product, which causes all of the consumers who search the entire market to switch to this firm. However, the searching population consists predominantly of low types, who will not upgrade to the profitable product. This places increased pressure on each sale of the high quality product to cross subsidise more loss-making sales of the base good, which cannot be sustained with a base good price below their competitors. As a consequence, firms shroud the price of higher quality items but advertise and partially compete in the market for low quality goods.

Inspired by Ellison's (2005) model, Gabaix and Laibson (2006) consider a related add-on sales mechanism. They model the idea that firms sell a basic product, such as a hotel room, and an add-on consumable, such as internet at the hotel. However, some consumers overlook the add-ons when committing to buy the base product. The key difference with Ellison's (2005) framework is that consumers in this model may buy both products, rather than one of the high or low quality items.

In Gabaix & Laibson's (2006) model there are also two types of consumers: Sophisticated and naive. Sophisticated consumers always foresee the need to purchase add-on products and take these into account when choosing the base good. However, naive consumers do not foresee the need to buy add-ons unless firms advertise them. When at least one firm advertises add-on products, some of the naive consumers become aware and act as if they were *sophisticates*. However, some naive consumers cannot be educated and remain unaware of add-on products.

They study a three period game in which duopolists decide whether to costlessly hide or advertise add-on prices alongside the price of the base product, in period 1. If firms shroud, sophisticated consumers are aware of the add-on but not its price and form Bayesian posteriors of the expected price. Naive consumers ignore the add-on completely. If at least one seller advertises the

add-on, the original proportion of sophisticated consumers plus some of the naive consumers are aware of add-ons and will take their price into account.

In period 2, consumers purchase a base good from one of the sellers and aware consumers have the option to exert the substitution effort to acquire the option of avoiding the add-on in the future.¹ All consumers who do not exert the substitution effort must purchase the add-on later. In the third period, the existence and price of the add-on is common knowledge to all buyers. Individuals who exerted effort in period 2 need only purchase the add-on if they wish. Individuals who did not exert the effort are obligated to buy the add-on.

Their main result is that for a wide range of parameter values it is individually profitable for firms to shroud add-on information. This leads to higher prices for add-ons and yields a social inefficiency in the form of the substitution effort that consumers who are aware of the add-on exert (Gabaix and Laibson, 2006). Moreover, in such a shrouding equilibrium it is not profitable for either of the firms to debias consumers due to the '*curse of debiasing*.' The nub of this effect can be explained using their simple example. Consider a consumer choosing between Hilton hotel and a Transparent rival, where Hilton charges a room rate of £35 and an add-on of £15. The market is perfectly competitive such that the supply cost is £50. When Hilton shrouds add-on prices, naive consumers are exploited. Now suppose that the Transparent hotel informs consumers of the add-ons and advertises an all inclusive price of £50. In this situation, the aware consumers continue to purchase the base good (hotel room) from Hilton but exert the substitution effort to avoid the add-on. This improves consumer welfare and reduces the revenue of Hilton because add-on sales decline. However, the informant rival does not benefit because consumers continue to buy from the firm on which their utility on the base product is highest. Therefore, in both Gabaix & Laibson (2006) and Ellison (2005), firms have an incentive to conceal some product information from consumers.

In many cases it is not clear which attribute of a multi-dimensional product consumers will find salient. For example, car insurance policies contain many dimensions and firms do not know which contingencies consumers will pay attention to. Different individuals may also choose different attributes. To capture this, Spiegler (2006) proposes a different style of multi-dimensional pricing to illustrate the idea that consumers pay attention to only some of the product's attributes but firms do not know which ones. More specifically, Spiegler (2006) assumes that consumers are boundedly rational in the sense that they cannot feasibly compare every contract on every dimension, which forces a short-hand sampling procedure: They select one attribute at random and buy from the firm that offers the lowest price on this dimension. However, the actual consumption

¹One natural interpretation is that consumers can download any files before arriving at the hotel to avoid using the internet.

utility is a random draw from the product's attributes because consumers do not know which contingency will emerge in the future.

This fits well with retail insurance. Drivers do not know if their car will be involved in an accident, stolen, suffer minor damage, break down or require replacing when choosing a policy. Similarly, home-owners do not know if they will be burgled, experience fire or water damage, or lose their house keys. Therefore, consumers compare policies using one of the contingencies but the contingency that occurs in the future determines the optimality of their policy.

In the two stage game, n firms commit to a pricing strategy that is captured by a cumulative distribution function (CDF). The CDF assigns individual prices to each attribute of the product. In the second stage, consumers randomly select one attribute and choose the firm that offers the best price on this dimension. Spiegel (2006) demonstrates that firms apply loss leader pricing strategies across attributes. Some attributes are priced competitively to attract buyers but other attributes are priced monopolistically to generate revenues from consumers who have purchased the product. In equilibrium, an increase in the number of sellers also leads to a mean preserving increase in the variance of prices. Therefore, increasing the number of sellers leads to greater price dispersion but average prices are unaffected. If consumers are risk averse, such that higher prices are over-weighted in their utility function, consumer welfare declines (Spiegel, 2006).

Ellison (2005), Gabaix & Laibson (2006) and Spiegel (2006) show how firms will exploit limitations in consumers' decision-making. However, a natural question is whether the results are preserved when firms are heterogeneous or other assumptions in these models are relaxed? Dahremoller (2013) and Wenzel (2014) explore two refinements of the Gabaix & Laibson (2006) structure and show that the incentives for firms to obfuscate are weaker when: (i) Firms vary in their costs, (ii) The proportion of naive consumers that become educated when a firm advertises add-on information is an increasing function of the quantity of unshrouding firms, rather than the fixed fraction of the unsophisticated consumers in Gabaix & Laibson (2006).

Dahremoller (2013) augments the model in three main directions. First, firms commit to an add-on disclosure strategy before setting prices. Secondly, a more specific demand function is imposed using a Hotelling line of product differentiation. These two refinements facilitate the third distinction: One firm incurs a positive cost of production for add-ons. Dahremoller (2013) demonstrates that in almost all cases there exists an incentive for at least one firm to unshroud add-on information. However, depending on the size of the fraction of consumers who are educated by the unshrouding policy, the transparent firm can be the efficient or the inefficient seller.

Dahremoller (2013) also demonstrates that unshrouding incentives are stronger if unshrouding informs consumers only of the price of the add-on at the unshrouding firm, rather than at all firms. In this sense unshrouding plays the role of informing buyers of the prices rather than making consumers aware of the existence of the add-on item.

Wenzel (2014) departs from Gabaix & Laibson (2006) by assuming a more elaborate unshrouding mechanism. The proportion of naive consumers that become sophisticated when one firm unshrouds is now an increasing function of the quantity of sellers that unshroud. Therefore, a natural interpretation for unshrouding in Wenzel's (2014) model is advertising: When more firms advertise add-on products, more consumers will recognise their importance.

Wenzel (2014) shows that the incentive to unshroud is now stronger, although in contrast with Dahremoller (2013) there continues to exist parameters for which shrouding is profitable. The incentive for a firm to unshroud now strengthens when a rival unshrouds because fewer consumers remain ignorant to add-ons. This produces a strategic complementarity between firms' strategies that is shut down in Gabaix & Laibson (2006). As a consequence, shrouding also becomes more appealing when the number of competitors increases.

In a further extension, Wenzel (2014) explores whether these results survive when advertising add-on information is costly for firms. In a similar way that Dahremoller (2013) resorts to a 'Hotelling street', Wenzel (2014) imposes greater structure on the demand function in the form of a Salop (1979) circular city model of product differentiation. Wenzel (2014) finds that when unshrouding costs are small, an increase in the number of sellers will continue to increase the likelihood of unshrouding. However, when unshrouding costs are large, unshrouding is non-monotonic with market concentration such that shrouding becomes more prevalent in less concentrated markets.

2.2.1.1 Endogenous Consumer Sophistication

In Gabaix & Laibson's (2006) workhorse model of add-on sales, some consumers are aware of add-ons whilst others are ignorant. Advertising the existence of add-ons also has a binary impact in the sense that the fraction of individuals who consider add-ons increases once, when at least one firm unshrouds (Gabaix and Laibson, 2006). However, it is not clear what fraction of individuals will be aware of the add-on products when firms do or do not advertise them. Therefore, a model in which the proportion of aware consumers is a continuous variable and endogenously determined can generate additional insights, which is the starting point for Carlin (2009).

Carlin (2009) considers an environment with many sellers of a single product. Firms simultaneously select both the price of their product and the complexity of the price, which is a scalar that aggregates to market complexity. The fraction of consumers who are able to evaluate prices is a decreasing function of market complexity. Whilst a scalar parameter for complexity offers many possible interpretations, it is reasonable to consider more complex prices as prices with more dimensions.

Price-complexity competition unfolds over two periods. In period 1, each firm $i \in n$ selects a price, $p_i \in [0, v]$ where v is the consumer's reservation price, and a level of complexity (k_i) for their

pricing structure, from the interval $[\underline{K}, \bar{K}]$. For every possible level of market complexity, some consumers (μ) understand prices whilst the remainder $(1 - \mu)$ do not, where $\mu : [\underline{K}, \bar{K}]^n \rightarrow (0, 1)$.

In the second period, consumers decide which seller to buy from. Informed consumers buy at the lowest price but uninformed consumers buy randomly. When a firm charges a high price, he will never attract informed buyers. This leads the the seller to also choose maximum complexity. The only incentive for a firm to deviate from maximum complexity is if the seller expects to attract informed consumers with a low price, in which case he prefers a larger fraction of informed buyers. Therefore, price and complexity are positively correlated and firms vary between maximum and minimum complexity. This implies that there exists a cutoff price, \hat{p} , such that when the firm draws a price above (below) \hat{p} , the firm applies maximum (minimum) complexity.

Firms also randomise their price but in a different way to Speigler's (2006) randomisation procedure. Speigler's (2006) randomisation is not a mixed strategy because firms assign single values to the price of each attribute. In Carlin (2009), firms randomise the price of a single attribute due to the classic trade-off between exploiting uninformed buyers and competing for informed buyers (Varian, 1980). Carlin (2009) also demonstrates that firms increase their complexity when the number of competitors increases, which can generate higher profits for firms.

Two simplifying assumptions underpin these results. Firstly, each consumer understands all of the prices or none of the prices in the market, which produces the extreme case of random or fully accurate purchasing decisions by buyers. Secondly, consumers who cannot understand prices, buy randomly. This implies that individuals cannot infer that more complex prices are more expensive. These assumptions are not unique to Carlin (2009) and underpin many classic models of price competition (Varian, 1980). However, relaxing these conditions can generate markedly different outcomes, as demonstrated in subsequent chapters of this thesis.

2.2.2 Price Framing Effects

When products feature multiple dimensions, it is intuitive that firms will manipulate the information disclosed to consumers. However, in many cases firms disclose all of the relevant information, yet consumers continue to experience difficulty when comparing competitors' prices. In this direction, firms may instead confuse consumers using complicated or individuated framing effects for their prices, which limits the ability of consumers to understand and compare prices. Natural examples include some firms using a single all-inclusive price when others use multi-part tariffs, and engineering unnecessarily complicated tariffs or product information.

The idea that the presentation of information can have a significant influence on decision processes is well grounded in psychology principles (Krishna et al., 2002).² Industrial economists, however, tend to focus less on the origins of consumer's biases and abstract from particular

²See Krishna et al. (2002) for a meta-analysis of price presentation effects in the marketing literature.

sources of confusion, to place greater emphasis on the resulting strategies for sellers. Therefore, complexity is often captured as a scalar parameter, where greater complexity enables fewer consumers to compare prices.

Salant and Rubinstein (2008) formalise an economic approach to frame-dependent choice by introducing framing as a secondary component to a standard choice procedure, in which individuals would otherwise select the utility maximising alternative. They develop an extended choice problem: (A, f) . X denotes the range of alternatives to which the choice function selects one element for each choice problem A . $f \in F$ denotes the framing of the choice problem, where F denotes the available set of frames. Frames include the order in which alternatives are presented or the existence of highlighted or default alternatives. Therefore, framing effects only influence choice because of psychological factors.

The extended choice function selects one alternative for every element in the pair (A, f) (Salant & Rubinstein, 2008). Salant & Rubinstein (2008) show that the extended choice description can capture information that is omitted by the standard choice approach and can model more elaborate decision-making processes of individuals. For example, in a standard choice model, an individual facing two alternatives, x and a , will choose x whenever: $u(x) \geq u(a)$, where $u(\cdot)$ denotes the utility function. However, suppose that product x is the default from which consumers are reluctant to switch. Switching only occurs if: $u(x) + \beta(x) \geq u(a)$, where $\beta(x) > 0$ captures the bias towards the default alternative (Salant and Rubinstein, 2008).

Eliaz and Spiegel (2011a) develop a model of frame-dependent choice in the context of competitive marketing. Consumers are randomly assigned to one of two firms, which captures their default supplier and buyers are reluctant to switch. Therefore, sellers must engage in costly marketing to persuade consumers to consider their product. Two firms simultaneously select a product (x) and a marketing strategy (M). Marketing determines whether a new consumer will consider their product. The continuum of identical buyers follow a two stage choice procedure. First, consumers define their consideration function, which is binary for every product-marketing pair in the sense that consumers either do or do not perceive the products to be substitutes. Second, if consumers accept the products as substitutes, they experience no difficulty in switching to a higher quality product. Consumers display inertia because they remain with their default provider if they do not perceive the rival's product to be a substitute or it is inferior.

Marketing cannot influence consumers' preferences, which are homogeneous and well-defined in terms of product quality. A fully rational consumer would consider all of the feasible options independently of the marketing that accompanies the product but consumers' aversion to switching creates friction. This ties to the idea of a consideration set in Spiegel (2006), except in that framework the consideration set comprises a subset of the attributes of each product, whereas the consideration set in this model comprises a subset of the products available.

Eliasz and Spiegler (2011a) abstract from prices as firms compete in market share minus costs, where the cost function contains two components. Production costs are positively related to product quality. Marketing costs are incurred to persuade a customer to consider the alternative product. Therefore, expected demand for a firm consists of half of the consumer population, plus any consumers that compare the product and find the product superior, minus any clients that compare the rival and choose to switch away.

Firms resort to extreme product differentiation. The highest quality product is advertised but the lowest quality product is not because only firms with higher quality products want to stimulate switching. Assuming costs are sufficiently small, there exists a symmetric equilibrium in which firms earn the max-min payoff. This is the same payoff that firms would achieve if buyers behaved rationally and exhibited no default bias. However, consumers are worse off than under rationality because some consumers buy inferior products. Eliasz and Spiegler (2011a) also show that their ‘persuade to consider’ framework elicits an efficiency in advertising because consumers who compare products always switch.

Piccione and Spiegler (2012) propose a related duopoly model in which the probability that consumers can compare prices is a direct function of firms’ framing decisions. Obfuscation occurs whenever firms select incomparable price frames. In contrast with Eliasz and Spiegler (2011a), framing prevents consumers from effectively comparing prices, rather than preventing consumers from identifying substitute products. Firms also maximise profit using prices rather than market share. They extend the traditional Bertrand model of competition to an environment in which consumers are limited in their ability to compare prices that are presented in different or complicated formats. Therefore, when buyers are impervious to framing, their model collapses to a standard Varian (1980) pricing game.

Two firms sell a single product to a unit demand consumer with a fixed maximum willingness to pay. The consumer is allocated to one firm at random to reflect prior purchasing decisions. He is subsequently able to compare his current contract with the contract offered by the rival with a certain probability, which is a function of firms’ framing decisions. If the consumer cannot compare the contracts or the alternative is inferior, no switching occurs. However, the consumer always switches to a comparable superior contract.

Their model revolves around a comparability structure, π , which determines the probability that the consumer can compare the prices of the two firms. If the two frames chosen are x and y , $\pi(x, y)$ denotes the probability that a consumer is able to compare the two contracts.³ Firms must choose a probability profile of price-format pairs that maximises expected profit. If there exists a universally comparable format, such that all consumers can perfectly compare any prices when

³Throughout the main analysis Piccione & Spiegler (2012) assume order independence such that $\pi(x, y) = \pi(y, x)$, which makes the default allocations irrelevant in determining whether the consumer can compare the prices.

at least one firm uses this format, both duopolists will adhere to the common format and earn zero profit in equilibrium. However, in the absence of a universally comparable format the results differ markedly.

The symmetric equilibrium hinges on a property of '*Weighted Regularity*': "An order independent comparability structure is weighted regular if there exists a probability distribution over the set of available formats such that, if one firm randomises over formats according to this distribution, the probability of a price comparison is independent of the other firms choice" (Piccione and Spiegler, 2012, p. 200).⁴ When weighted regularity is satisfied, all firms earn the max-min payoff. Furthermore, measures to increase comparability, in the sense that $\pi(x, y)$ increases for every pair of framing decisions x and y , will improve consumer welfare, reduce firms' payoffs and stimulate switching. However, when weighted regularity is violated this need not be the case and increasing the comparability of product frames can lead to higher prices and less switching.

Chioveanu and Zhou (2013) independently propose a model of competition in which consumers can become confused by firms adopting different product frames. When there are only two firms the market, their price-frame model becomes a special case of Piccione & Spiegler (2012). However, they focus in great detail on distinguishing between frame complexity and individuated framing as separate but interactive sources of consumer confusion. In addition, they extend their model to an oligopoly market with many firms to understand the impact of competition. Firms follow the same timing as Piccione & Spiegler (2012) by choosing between prices and price frames simultaneously and selling to a unit mass of potential customers.

In the two firm case, Chioveanu & Zhou (2013) show that when consumers are most confused by firms choosing different frames, rather than an intrinsically complex frame, pricing is independent of the choice of frame and profit is increasing in the proportion of confused consumers. In contrast, when consumers are most confused by firms choosing an intrinsically complex frame, firms that select a more complex frame charge higher prices. Both distributions generate positive profits.

In the oligopoly environment they demonstrate the importance of distinguishing between the two sources of consumer confusion and consumers adopt a two-stage dominance based choice rule. Firstly, all offers that are comparable to another and dominated are discounted. Secondly, consumers buy from a randomly selected firm from the remaining two. They also characterise the equilibrium in terms of a more general purchasing rule than pure randomisation by allowing consumers to prefer the simple or complex offer but they do not analyse the implications of such preference behaviour and restore the random purchase rule throughout their analysis. I develop this idea further within this thesis.

⁴More simply, weighted regularity states that if there are two available formats for pricing a product (such as miles per litre and miles per gallon for fuel efficiency of a vehicle), then if firm 1 randomises equally between the two frames, the probability of a comparison by the consumer is 0.5, independent of the rival's decision.

Chioveanu & Zhou (2013) show that when frame differentiation is more confusing than frame complexity, prices and price frames are no longer independent with many firms and the price distributions for the two frames cannot be placed in an ordinal ranking. An increase in competition leads firms to increase the probability placed on the complex frame because adopting a different frame becomes more difficult, which can generate a negative correlation between profit and concentration. In contrast, when consumers are most impaired by complex framing, the higher price distribution is associated with the more complex frame in the oligopoly model and firms adopt the complex frame with greater probability as competition intensifies. Industry profit can also increase with the number of competitors.

In the preceding discussion, framing effects are permanent: A consumer can or cannot compare products when making a purchase. However, in many environments framing effects are transient in the sense that they wear off at a later date. For example, the value of a designer dress or jacket may seem higher for the consumer in the context of the brightly lit studio than when the individual returns home (Salant and Siegel, 2015). Salant and Siegel (2015) propose a model in which framing continues to be utility irrelevant for buyers but with three main differences. Firstly, firms sell a product pair (x, t) , where x denotes quality, t denotes price and production cost ($c(x)$) is increasing in quality. Secondly, there exists a single monopolistic supplier. Thirdly, framing temporarily influences consumers' valuations of products.

Consumers buy as follows. In stage one, each buyer selects a product that maximises their *perceived* utility (u^f), which framing can inflate. In stage two, framing wears off and the consumer re-evaluates their purchase according to their true valuations, (u). If the *real* utility exceeds the utility associated with the outside option of making no purchases, the consumer retains the product. If not, the product is returned for a full refund. The motivation for firms to engage in framing is, therefore, different from Eliaz and Spiegler (2011a), Piccione and Spiegler (2012) and Chioveanu and Zhou (2013) as sellers seek to inflate consumers' valuations of products, rather than limiting comparability.

This maximisation and verification choice procedure connects to the more general idea of reservation prices, according to which a consumer is willing to pay no more than their reservation threshold, which firms usually cannot manipulate and do not price above. However, in Salant & Siegel (2015), consumers may overpay for the product in the first stage but these products are always returned.

They explore two main questions. First, how does transient framing interact with regulation policies that compel firms to offer a basic product, in addition to more expensive and higher quality items? Second, do framing strategies always produce higher profits when consumers are heterogeneous? The strategy for firms is to design a menu of products (quality-price combinations), accompanied by a frame, that maximises profit.

Consider first the case of a market in which regulators impose a basic product: (\bar{x}, \bar{t}) and consumers hold identical preferences. In the absence of framing, such a policy places a lower bound on the minimum utility a consumer can obtain, which limits the prices that sellers can charge for higher quality products.⁵ Therefore, consumers obtain a high quality product at a lower price than would arise without the regulation. However, firms can circumvent this restriction by manipulating consumers' valuations of the products. Providing that there exists a frame that ensures consumers prefer the socially optimal higher quality item to the basic product, firms can charge higher prices. If such a frame does not exist, the optimal contract involves firms investing in excessively high quality products and pricing at the socially efficient level. This leads to higher profits and lower social welfare than the optimal contract without framing. Most interestingly, framing reverses the intended consequences of the 'basic contract' policy, as surplus shifts back towards the seller.

In a market with heterogeneous buyers, however, framing does not necessarily increase profit. Consider a market with two types of consumers: Low and High, where Low types are always willing to pay less for a marginal increase in quality, independently of framing. The firm can offer separating contracts to the two types of consumers (Salant and Siegel, 2015). Framing has two effects on profit. Firstly, if a consumer is willing to pay t for product x without framing, the consumer will be willing to pay more for the same product with framing, which allows the firm to increase prices. This follows the case of the regulated market. However, a secondary effect arises because the same product x now appears less attractive to consumers, relative to the higher quality contracts on the menu. However, these higher quality products may be offered at a price that is above the consumer's reservation threshold. Therefore, contracts designed for high-type buyers may cause low-types to leave the market without purchasing and the seller must reduce the price of the basic good to ensure participation. Therefore, framing can reduce profit (Salant and Siegel, 2015).

In an extension, Salant & Siegel (2015) analyse the role of 'highlighted options' that are more prominent on the menu. They consider a monopolistic market for insurance with high and low risk individuals. The highlighted contract replaces the outside option of not purchasing any of the contracts because consumers are loss averse and experience regret if inferior products are purchased, thereby distorting the consumer's reference point. They show that the firm will highlight the premium product and may choose not to offer a contract to high risk individuals.

In a related paper, Spiegler (2014) generalises Piccione & Spiegler's (2012) framework to extend their notion of weighted regularity and highlight two key points. First, weighted regularity is satisfied in many competitive environments, which facilitates tractable analysis of the interaction

⁵A consumer will only purchase the higher quality item if: $u(\bar{x}) - \bar{t} \leq u(x') - t'$, where (x', t') is the higher quality product.

between framing and the number of competitors. Second, many multi-dimensional promotion problems can be recast as specific forms of framing.

In the model, two firms compete by choosing a product-marketing pair. The product, a , denotes the alternative that the firm chooses to supply, which can be interpreted in quality terms. $M(a)$ denotes the marketing messages available for alternative a . The combination of firms' marketing decisions is represented by a symmetric function, $\pi : MXM$, which maps from every combination of marketing messages to a probability profile over the possible frames, $f \in F$. Formally, $\pi(M, M) \rightarrow \Delta F$ such that $\pi_f(m_1, m_2)$ is the probability that the consumer adopts frame f when appraising the products. The function $p : A \rightarrow \mathbb{R}$ determines the profit for the firm from selling each alternative. That is, $p(a)$ is the payoff from each unit of sale of good a .

The final ingredient is a probabilistic choice function: $s_i(a_1, a_2, f) \in [0, 1]$. This defines the probability that the consumer will purchase from firm i , conditional on the range of alternatives $\{a_1, a_2\}$ and the frame. That is, s_1 is the probability that the consumer buys from firm 1 and $s_1 + s_2 = 1$ such that consumers always purchase one of the products. The firm's objective function is:

$$p(a_i) \cdot \left[\sum_f \pi_f(m_1, m_2, f) \cdot s_i(a_1, a_2, f) \right] \quad (2.1)$$

To capture the elegance of this framework I report two applications provided by Spiegel (2014): Gabaix & Laibson's (2006) add-on model and Piccione & Spiegel's (2012) framework of limited comparability. Recall the idea of Gabaix & Laibson (2006): Firms supply a base good and a more profitable add-on but only some consumers pay attention to the add-on. If at least one seller unshrouds, some of the naive consumers become educated and take the add-on into account.

The alternatives are a vector of the two products: (a^1, a^2) , that specify the quality of the two characteristics of the overall product. The profit per sale is equal to the consumers' reservation price minus any investment in quality: $p(a^1, a^2) = 1 - (a^1 + a^2)$. A consumer may experience shrouded or unshrouded add-ons. Let $F = M(\cdot) = \{0, 1\}$, where $m = 1$ corresponds to shrouding and $m = 0$ corresponds to unshrouding. The consumer becomes aware of the add-on only if the frame is $f = 1$, otherwise the add-on is neglected. Spiegel (2014) imposes $\pi_1(m_1, m_2) = m_1 m_2$. Therefore if at least one firm unshrouds ($m = 0$), the consumer foresees the add-on with certainty. The choice function is $a_i^1 + (1 - f)a_i^2$, where f is binary and captures the two states.

A similar game description can be provided for Piccione & Spiegel (2012). Consumers are initially divided equally across sellers and framing decisions influence whether the consumer is able to compare the rival's product. The alternatives are priced according to $p \in [0, 1]$ and the

set of possible frames is binary. $f = 0$ defines incomparable formats and $f = 1$ defines perfectly comparable formats. The choice function becomes:

$$s_1(p_1, p_2, f) = \frac{1}{2} + f \cdot \text{sign}(p_2 - p_1) \quad (2.2)$$

When formats are incomparable, each firm sells only to its default population. However, when frames are comparable, all consumers buy the cheapest product. Therefore, multidimensional pricing mechanisms and models of competitive framing are intimately connected concepts that can be captured by specifying the corresponding frame, available marketing messages, consumer choice rule and the firm's objective function. Spiegler (2014) also demonstrates that government regulations that impose a single frame on firms, need not maximise consumer welfare. The intuition is that if firms instead selected a different common frame, competition may be stronger.⁶

2.2.3 Saliency

Saliency is an intimately related idea to multi-dimensional sales and strategic framing in which firms emphasise some features of their product but conceal others (Kim and Kachersky, 2006). However, in Industrial Economics, the relative saliency of each attribute of a product is determined by the value of the attribute, relative to the value of the same attribute of other products. That is, an attribute becomes salient when its value differs most from the average value in the market for that attribute.

A formal economic approach to saliency emerged only recently, most notably from Bordalo et al. (2013).⁷ They propose a model in which products contain multiple attributes but: (i) Individuals are aware of and understand all of the attributes, (ii) The relative weight placed on each attribute varies across contexts for the same individual. In contrast with Salant & Siegel (2015), consumers correctly value the products' quality and firms cannot inflate or conceal either of the quality or price dimensions. They take rational decision making as a starting point and introduce the idea that individuals are 'salient-thinkers' in the sense that greater attention is allocated to 'surprising' product attributes. These are attributes that differ most from the average value.

More specifically, in a market with N products, where each product consists of a quality-price pair that are both expressed in monetary units, a rational consumer will purchase the product that

⁶For example, a graph showing the prices of one firm against a competitor may have axes for the price ranging from 0 to £5000. If all firms adopt the same graph and scale, consumers can reasonably identify the cheapest seller. However, if all prices lie between £4500 – £5000, a graph indexed from zero to the maximum price will understate the absolute price differences. Deviating from the default presentation format to a graph scaled to the relevant pricing range can highlight the dispersion in prices and stimulate consumer engagement (Spiegler, 2014).

⁷See Kim and Kachersky (2006), and the references contained within, for a conceptual framework and review of saliency effects in the marketing psychology literature.

maximises his objective utility. That is, consumer i values product $k \in N$ as their valuation of the quality of the product minus its price:

$$u_i = q_k - p_k \quad (2.3)$$

The rational thinker weights price and quality equally but a salient thinker allocates less attention to one of the features. If quality is salient, prices are under-weighted in the utility function:

$$u_i = q_k - \delta p_k \quad , \quad \delta \in [0, 1] \quad (2.4)$$

If prices are salient, the importance of quality is diluted:

$$u_i = \delta q_k - p_k \quad , \quad \delta \in [0, 1] \quad (2.5)$$

δ captures the degree of inattentiveness to the less salient attribute and is an exogenously determined fixed value throughout their analysis.

The next ingredient to their story is the endogenous determination of which attribute is salient according to the average value of each attribute in the market and the dispersion of each products' attributes from the average value. This distinguishes the idea further from Gabaix & Laibson (2006) because in this model consumers are fully informed. When the differences in quality across products are large, relative to the dispersion in prices, quality becomes a focal point for consumers. When the relative price differences are larger than the relative quality differences, prices become salient. To explain the main intuitions I adapt their simple example (Bordalo et al., 2013).

Consider a consumer choosing between a low and a high quality bottle of wine at a supermarket, priced at £10 and £20 respectively. The individual believes that the expensive wine is 1.5 times the quality of the cheaper wine. At supermarket prices, the consumer is only willing to pay £15 for the upgrade and the cheaper bottle is purchased. The following week, the consumer faces the same decision at a restaurant but the prices are equally marked up to £50 and £60 respectively. If the consumer must pay £50 for the low quality product, a £10 upgrade is now relatively cheap for a bottle of 1.5 times higher quality and the consumer upgrades.

To formalise these arguments, consider a market with N products. Define the average quality and price as $\bar{q} = \frac{\sum_k q_k}{n}$ and $\bar{p} = \frac{\sum_k p_k}{n}$, respectively. This will determine the offers that a consumer perceives to be a good deal. Let $\sigma(a_k, \bar{a})$ denote the salience function for each attribute and product k . The salience function determines which of the attributes consumers pay greatest attention towards and follows the properties of ordering and diminishing sensitivity.⁸

These conditions imply that quality is salient if and only if $\sigma(q_k, \bar{q}) > \sigma(p_k, \bar{p})$ (Bordalo et al., 2013), which in the two product case translates to:

⁸Definition 1: Ordering: Let $\mu = \text{sgn}(a_k - \bar{a})$. Then for any $\epsilon, \epsilon' \geq 0$ with $\epsilon + \epsilon' > 0$, we have: $\sigma(a_k + \mu\epsilon, \bar{a} - \mu\epsilon') > \sigma(a_k, \bar{a})$

Definition 2: Diminishing Sensitivity: For any $a_k, \bar{a} \geq 0$ and all $\epsilon > 0$, we have: $\sigma(a_k + \epsilon, \bar{a} + \epsilon) < \sigma(a_k, \bar{a})$ (Bordalo et al., 2013).

$$\frac{q_h}{p_h} > \frac{\bar{q}}{\bar{p}} \quad (2.6)$$

(q_h, p_h) denotes the attributes of the higher quality product. Intuitively, products with a higher quality/price ratio are preferred by consumers. Bordalo et al (2013) demonstrate that their framework can capture several inconsistencies in choice patterns.

Two main applications are relevant for industrial economists. Firstly, the reference point, given by the average value of each attribute, makes individuals susceptible to firms' tampering with this value. This may involve introducing a decoy product, which is inferior to all other contracts. In the wine example at the supermarket, $\bar{p} = 15$, $\bar{q} = 20$, $q_h = 30$ and $p_h = 20$. Therefore: $\frac{q_h}{p_h} < \frac{\bar{q}}{\bar{p}}$ such that price is salient and the cheaper good is purchased. However, introducing a decoy: $(30, 30)$, distorts the consumer's reference point: $\bar{p} = 20$, $\bar{q} = 26.66$, $q_h = 30$ and $p_h = 20$. Therefore, $\frac{q_h}{p_h} = 1.5 > 1.33 = \frac{\bar{q}}{\bar{p}}$, such that quality is salient and the consumer upgrades.

The second application is misleading sales. Consumers form reference prices based on the expected price in the store but the expected price is a weighted average of both the sale and non-sale prices. Therefore, by inflating the non-sale price, firms can increase the price a consumer expects to find, which exaggerates the price saving because consumers over-weight prices. Therefore, when products are only sold by a single supplier, the firm can shift demand towards higher quality goods. However, when many firms sell the same product, quality will naturally be under-weighted. As a consequence, buyers reduce their reservation price, which counteracts firms' incentives to engage in misleading sales.

In a companion article, Bordalo et al. (2016) introduce competition and analyse the impact of asymmetric costs across sellers and drastic changes in the cost of production. Two retailers, $k = \{1, 2\}$, produce a single good of variable quality (q_k) with a cost function: $c_k(q_k) = F_k + V_k(q)$. F_k captures the cost of producing each unit of the product and $V_k(q)$ captures the cost of increasing quality. The cost function is increasing in quality, convex and satisfies $V_k(0) = 0$. The firms sell to a population of identical buyers who each buy one product from the choice set: $C = \{(q_1, p_1), (q_2, p_2)\}$, where (q_k, p_k) denotes the quality-price bundle of firm k . Therefore, one firm supplies the entire market.

The game develops over two periods. In the first stage, firms commit to supplying a particular quality of product ($q_k \in [0, \infty)$). In the second stage, quality commitments become common knowledge and prices are selected. The idea of salience now revolves around an 'attention externality' which operates as follows. If a firm selects a high level of quality, the product becomes more attractive to buyers for two reasons. Firstly, consumers prefer higher quality goods. Second, an increase in quality tilts the quality-price ratio towards a quality-salient equilibrium, which highlights the lower quality of the rival (Bordalo et al., 2016). Similarly, a low price is intrinsically

preferred by consumers but also creates an externality by highlighting the high price of the rival. They demonstrate that the cost of production determines whether quality or price salience will emerge.

Bordalo et al (2016) show that salience effects can generate supra-competitive profits. When quality is salient, consumers overvalue the high quality product, which allows the higher quality seller to increase prices. When price is salient, consumers inflate the value of the low price product, which allows the low quality seller to increase prices. When setting quality, the relative values of marginal and average cost are critical. In the symmetric case, where neither firm benefits from a cost advantage, sellers follow the same strategy and salience effects are suppressed. However, in contrast to the rational benchmark, firms will select an inefficient quality level when consumers are salient-thinkers for the following reason. If quality exceeds the efficient level, a deviation by one seller to a lower (socially optimal) quality-price pair is not profitable because quality becomes salient and consumers undervalue the price reduction. If quality is below the efficient marker, an increase in the quality and price is not profitable because the consumer overweights the increase in price. This makes the product appear less attractive.

Quality provision is also weakly increasing in the unit cost component: F_k , due to the interplay between marginal and average costs. Prices also increase with F_k . Therefore, consumers become desensitised to price differences in markets where the unit cost of production, which is independent of quality, is high. Bordalo et al (2016) argue that this novel property can explain the observed quality competition in markets for luxury goods and price competition in markets for basic goods.

Shocks to the unit cost influence investments in quality significantly. The impact of reducing the marginal cost of quality for one seller depends on the magnitude of the cost advantage and the relative value of the marginal and unit costs. When the cost advantage is large, or if the unit cost is large, the more efficient seller monopolises the market by upgrading his product. That is, investing in higher quality and charging a higher price. The increase in quality draws consumers' attention away from prices because average prices are already high due to the large unit cost, and the cost saving can shift the market from price to quality salient. In contrast, when the cost advantage is small and unit costs are low, such that average prices are already low, consumers will focus on prices. Therefore, additional investments in quality are under-valued by consumers and the optimal strategy involves no change to the product.

These dynamics are similar in spirit to add-on pricing because in both cases firms reduce the value of dimensions that consumers pay less attention to. However, in add-on pricing games, some consumers are naive and neglect add-ons. In salience frameworks, consumers are fully aware of the attributes of all products but choose to weight them differently.

A common theme throughout these studies is a focus on partial equilibrium. That is, we study one market in isolation to cast light on specific comparative statics or market interactions. However, in practice, consumers purchase many independent products from several markets and they cannot analyse every price in every market. Rational individuals will, therefore, allocate their attention towards markets where the expected price saving is largest, which is the starting point for De Clippel et al. (2014). They show that this overarching constraint for consumers introduces a new layer of cross-market competition because firms charge the highest price possible without drawing attention to their market.

In their model, there exists M independent markets, where each market contains two independent sellers. Each firm operates only in one market. Each consumer has a reservation price of 1 for the product in each market and is initially allocated to the market leader. The buyer must decide in which markets to learn the price of the less prominent seller but because attention is a scarce resource, the consumer can only investigate a subset of the markets from which they will purchase. Therefore, the consumer is destined to buy from the market leader unless he invests in learning the alternative prices. This creates inertia, similar to Eliaz and Spiegler (2011a).

The consumer buys the bundle that maximises his utility, subject to the attention constraint. In the absence of attention constraints, the model collapses to a game with complete and perfect information. They address two questions: How should consumers allocate their attention? What is the impact for firms' strategies and consumers' surplus? Consumers vary only in their total attention constraint. α_k denotes the proportion of consumers that are able to examine at most $k \in M$ markets. $\alpha_M = 1(0)$ defines complete (in)attention.

Write the utility for the consumer from purchasing products $(x_1, x_2, \dots, x_M) \in \{0, 1\}^M$ at prices (p_1, p_2, \dots, p_m) as:

$$\sum_{m=1}^M (1 - p_m) x_m \quad (2.7)$$

In period 1, both prominent and non-prominent firms independently set the price for their good to maximise profit. In period 2, consumers observe all of the leaders' prices and choose which markets to investigate further. De Clippel et al (2014) study a perfect Bayesian equilibrium and they restrict their focus to partially symmetric strategies such that every prominent firm and every non-prominent follows the same strategy, although the strategies between the two firm types will differ.

De Clippel et al (2014) demonstrate that partially attentive consumers introduce a novel dimension of inter-market competition and market leaders charge lower prices than if consumers

were fully attentive. By charging a higher price, the seller generates additional profit from consumers who buy from the firm but the higher price also increases the probability that attention constrained individuals will investigate that market. The less prominent firm must always set a lower price than the prominent firm to achieve any sales. Therefore, any consumer who investigates the market will switch, which connects to the efficiency in advertising of Eliaz and Spiegelger (2011a).

An increase in the fraction of captive consumers (α_0), who always buy the prominent product, will also lead to lower consumer surplus because both firms increase their prices. However, increasing the fraction of fully attentive buyers (α_m) does not benefit consumers. Instead, consumer surplus increases when the proportion of fully attentive consumers declines because the incentive for the prominent firm to 'hide' by reducing its price, strengthens. Therefore, consumer surplus is maximised when consumers search either zero or one market, such that the distribution of attention satisfies:

$$(\alpha_0, \alpha_1, \dots, \alpha_m) = (\alpha_0, 1 - \alpha_0, 0, 0, \dots, 0) \quad (2.8)$$

In contrast, consumer surplus is minimised when the cross-market effects are weakest such that all consumers are fully attentive or inattentive:

$$(\alpha_0, \alpha_1, \dots, \alpha_m) = (\alpha_0, 0, 0, \dots, 1 - \alpha_0) \quad (2.9)$$

2.2.3.1 Attention Grabbing Products

When firms select which products to stock, the expected profit from the sale of each item only provides part of the picture. Rather, the total value of offering a product is determined by both the direct sales of the item and the indirect demand it generates for other products in the store. Therefore, it may be optimal to stock items that do not cover their cost or are never intended for sale, to stimulate sales of other products (Eliaz and Spiegelger, 2011b).

Eliaz and Spiegelger (2011b) develop a model in which consumers are initially distributed evenly between two firms and buyers are limited in their attention. Therefore, unless a firm is able to attract consumers' attention to their menu, they will never consider any of its products. This creates an opportunity for firms to add attention grabbers to their menu, which are utility-irrelevant for consumers and costly to make available, but entice consumers to the menu. They consider the optimal use of attention grabbers and the implications for sellers' payoffs and consumer switching.

For example, consider a consumer travelling to his favourite restaurant who, due to limited attention, does not usually consider the menu of the restaurant next door. If the competitor introduces one attention grabber, the consumer is intrigued and inspects the menu of the rival. He

may not buy the attention grabber, but it acts as a door-opener (Eliaz and Spiegel, 2011b). Alternatively, consider a television network. The distributor may promote a programme that captures consumers' attention to the channel, even though the consumer will not watch this particular programme (Eliaz and Spiegel, 2011b).

In the game, two firms maximise their market share by designing an optimal menu of products (M) from the products available (X). They must attract new consumers and retain their default clientele, whilst minimising the cost of offering their menu. The cost of offering a menu M is the sum of the cost of offering each item listed: $c(M) = \sum_{x \in M} c_x$, where c_x is the cost of listing item x . This captures the trade-off between holding an attention grabbing inventory and minimising operating costs.

Consumers are identical and follow a two step choice procedure, which is related but distinct to Eliaz & Spiegel (2011a). In stage 1, consumers decide whether to consider the menu of the rival seller according to an attention function, which is written:

$$f : X \cdot P(X) \rightarrow \{0, 1\} \quad (2.10)$$

$P(X)$ denotes the list of available menus. The default supplier cannot influence whether the consumer considers the rival. This differs from price-frame competition in which consumer sophistication is jointly determined by the firm's choices (Carlin, 2009; Piccione and Spiegel, 2012).

In stage 2, consumers choose the superior alternative from the menus considered. The consumer will only switch if there is an item on the rival's menu that captures their attention and there is subsequently an item on that menu that the individual prefers. If attention is infinite, consumers consider every menu and sellers provide the best menu possible, whilst minimising costs. Therefore, attention grabbers will not be offered, firms earn the max-min payoff and consumers never switch. However, when attention is scarce, attention grabbers play a pivotal role.

Eliaz and Spiegel (2011b) consider two classes of attention functions: Ordered and similarity based. Ordered assumes that products can be ranked according to their intrinsic attention attracting value because some products attract more attention than others. In this case a consumer will only consider the rival's menu if it generates more attention than their default menu. Similarity-based assumes that consumers seek familiarity in their consumption decisions and will only purchase a product that is sufficiently similar to their default item. Under both types of attention functions, firms earn the same max-min payoff that they would achieve if consumers were fully attentive as the cost of attention grabbers is exactly offset by savings from providing lower quality menus. However, the equilibrium itself is very different.

Under ordering, firms select a probability distribution over the available menus, which includes positive weight on both optimal and inferior menus. Firms list at most one attention grab-

ber and only when the firm is offering the consumers' most preferred product. The probability placed on sub-optimal menus is also increasing in the cost of providing the best attention grabber.

Under similarity, Eliaz and Spiegel (2011b) impose max-max preferences for consumers, such that firms only offer menus that contain a single content item. The two products are located on the real line and their proximity reflects the degree of similarity. They find that firms vary between consumers' preferred menu and inferior menus, and the probability placed on the most preferred menu is bounded from below by a value that declines with the cost of supplying sub-optimal products.

In both cases, limited attention generates positive consumer switching, which generates a social loss. This is somewhat akin to the deadweight loss in Gabaix & Laibson (2006). In that market, some consumers foresee high prices for add-ons and exert socially wasteful substitution effort to avoid them. One difference, however, is that consumers in Eliaz and Spiegel (2011b) are identical. In an extension, Eliaz and Spiegel (2011b) introduce heterogeneous consumer preferences. They prove the existence of a symmetric solution that resembles the benchmark with identical buyers and show that whenever a menu features an attention grabber, consumers who are not already consuming from that menu will choose to switch. Therefore, attention grabbers alert consumers to a superior alternative, which connects to the previously identified 'effective marketing property' (Eliaz and Spiegel, 2011a).

2.2.4 Search Costs

Classic models of price competition assume that some consumers can evaluate every price available but, due to cognitive constraints, others cannot evaluate any (Varian, 1980). An alternative methodology is to introduce search costs, which include the transport cost of visiting the store, the mental effort of concentrating on prices or the time required to evaluate prices (Tirole, 1988). Search costs rationalise an individual's decision not to acquire market information because the marginal cost of search exceeds the anticipated gains. Therefore, heterogeneity amongst consumers stems from divergences in their personal '*cost of search*' rather than their cognitive ability or attentiveness.

Obfuscation in search models involves sellers inflating the cost of search. This reduces the amount of information that consumers acquire. Examples include a firm hiding the final price of a product on a website or locating cheaper products in obscure locations in a store. Search-based obfuscation has attracted considerable attention in recent years as the advent of new information technologies and price comparison websites should streamline the comparison process for consumers and minimise search costs (Bakos, 1997; Ellison and Ellison, 2009).

Wilson (2010) develops a model in which two firms choose the *search cost* associated with the consumer learning its price, denoted s_i for each firm i . In the second stage, firms select their

prices conditional on the observable search costs in the market. In the third stage, consumers decide how many of the duopolists to learn the price of and the order in which to search the firms, before purchasing. Consumers have a common valuation for the product but they vary in their cost of search. An exogenously determined fraction of the consumers incur a positive search cost, whilst the remainder search for free. Therefore, increasing the cost of search reduces the amount of search that costly search consumers undertake.

Wilson (2010) demonstrates that the two firms select asymmetric search costs, which induces costly search consumers to consider the firm with the lowest search cost first. This furnishes the low search cost firm with market power over the costly search consumers. As a result, the low search cost firm prices its product in a way that guarantees costly search consumers choose to buy rather than searching again. The firm with the higher search cost will never see any of these potential customers and will offer a marginally lower price to attract costless search buyers. Obfuscation therefore enables both sellers to earn higher payoffs. One assumption underpinning Wilson's (2010) story is that consumers know the cost of searching each firm *ex ante* because consumers search firms in ascending order of their search cost. When this assumption is violated, the mechanism fails.

Ellison and Wolitzky (2012) develop two models based on the underlying framework of Stahl (1989) that, amongst other refinements, removes the assumption that consumers know the time required to search each firm *ex ante*. Instead, consumers know only the distribution of the time required to learn the prices of firms. Consumers have identical demand functions, $D(p)$, that decrease with price (p) and consumers vary in their cost of search.

In Stahl (1989), consumers are divided into two types: Costly and costless searchers. Costly searchers incur a fixed per unit cost for every unit of time (t) that is expended during the search process but costless searchers do not. The proportion of costly searchers and search cost are given exogenously and firms only choose prices. This leads sellers to randomise their price in a similar intuition to Varian's (1980) model of sales.

Ellison & Wolitzky (2012) develop Stahl's (1989) framework by allowing firms to select the time required for a consumer to learn their price. Therefore, searching firm i requires $\tau + t_i$ time, where τ is exogenous and common across firms and $t_i \geq 0$ is the endogenous additional time required to learn the price of firm i . Firms simultaneously select the price and search time, which differs from the sequential timing structure in Wilson (2010). Consumers subsequently make their search and purchasing decisions. Costless searchers learn all of the prices and choose the cheapest firm. Costly searchers choose one firm randomly to begin their search, which incurs the first unit of the search cost. The costly searcher subsequently decides whether to search again, conditional on the price offered by the first firm, the anticipated savings of further search and the expected cost of further search.

In their first model, Ellison and Wolitzky (2012) introduce a strictly convex disutility function of search: $g(t)$. They characterise the symmetric sequential equilibria of the game, concentrating on equilibria in which costly searchers are incentivised to engage with the market and buy. They find that many of the properties of Stahl (1989) survive in the environment with convex disutility, conditional on obfuscation being costless for firms. Sellers price from non-atomic distributions but the distributions are always bounded from above at the monopoly price, which is higher than the obfuscation-free level. However, to sustain this upper bound a minimum level of search friction is required. This places a lower bound on the search time that firms choose. Consumers are therefore worse off because search costs and prices are both higher (Ellison and Wolitzky, 2012).

In an extension, obfuscation becomes costly for sellers. This pins down the optimal level of obfuscation to the minimum amount of time that ensures that even when a costly search individual is offered the maximum possible price in the distribution, they will buy it, even though they are certain that the price elsewhere will be cheaper. Firms with a higher price choose a higher search time, whilst a firm with a low price may not obfuscate at all. This echoes Gabaix & Laibson (2006) in the sense that hidden prices are most expensive. The prediction also contrasts with the results of Wilson (2010) in which the firm with the lowest obfuscation charges the highest price. The divergence is understood by considering the motivations for firms to obfuscate. In Wilson's (2010) set-up, obfuscation directs costly searchers towards the rival and causes the rival to increase his price, which allows the obfuscating firm to increase his price. In Ellison and Wolitzky (2012), however, obfuscating firms benefit directly from the additional time they exhaust from consumers.

Ellison and Wolitzky (2012) also show that when obfuscation is costless, any reduction in the natural length of time associated with learning a firm's price, τ , will be matched by an increase in obfuscation by firms. This connects with the empirical discussion later in this review where sellers manufacture search frictions when technological advances reduce natural impediments to search (Bakos, 1997; Ellison and Ellison, 2009). With costly obfuscation, however, this may not be the case because the cost may outweigh the benefit from suppressing search and it is possible that consumers may visit two stores before making their purchase. More surprisingly, when obfuscation is costly, obfuscation is non-monotonically related to the price that a firm charges. However, several markers indicate that prices will continue to be positively related to obfuscation in most cases.

In their second model, Ellison & Wolitzky (2012) revert to a linear disutility function for consumer search and costless obfuscation for sellers but introduce common uncertainty, for both firms and consumers, regarding the value of the exogenous element of the search time, τ . Therefore, when firms select their obfuscation, they do not know the total amount of time that consumers incur to learn their price and consumers cannot decompose the total search time they incur at a

particular store into firm-induced and exogenous components. This distorts consumers' expectations of the cost of further search.

Specifically, suppose that the disutility of searching for time t follows: $g(t) = t$ and let θ denote the distribution of possible values for the exogenous search time, τ . $\mathbb{E}[\theta] = 0$ and θ is distributed with a continuous density $h(\theta)$ over $[\underline{\theta}, \infty)$. The parameter also satisfies $\underline{\theta} > -\tau$ and is drawn once at the beginning of the game (Ellison and Wolitzky, 2012). They focus on the minimum obfuscation level associated with each price. That is, for every price (p) there exists a time ($t^*(p)$) such that profit is maximised at the minimum possible search time, which introduces greater determinacy into the equilibrium. Once buyers learn \tilde{t} and the price, they know that: $\theta = \hat{\theta} \equiv \tilde{t} - (\tau + t^*(p))$.

Ellison and Wolitzky (2012) show that costly searchers never search twice, replicating the search patterns of the first model. However, there are many new equilibria that include price distributions bounded from above at a price below the monopoly level. This model also discloses an excess obfuscation problem in which obfuscation harms sellers and buyers. Sellers must ensure that costly searchers do not search twice, even if the exogenous part of the search cost is drawn to be the lowest possible value (Ellison and Wolitzky, 2012). Therefore, firms choose a higher level of obfuscation than is necessary and prices also decrease as this excess obfuscation problem grows. However, this does not benefit consumers because the additional dis-utility from search negates the price reduction.

2.2.4.1 Multi-Product Retailing

Search costs play a more natural role when consumers buy several different products. Shelegia (2012) analyses multi-product pricing and demonstrates that the co-movement in prices between products depends critically on whether the goods are economic complements or substitutes. The explanation follows that the economic status of products affects the price that consumers are willing to pay for the combination of products, relative to their individual valuations.

In Shelegia's (2012) model, however, the consumer's search process is exogenous: Some consumers search both stores (*shoppers*) whilst others only visit one (*captives*). The two sellers each supply, without cost, two identical products (a and b) to the unit mass of consumers. Consumers decide whether to buy either or both of the products but fraction θ visit only one randomly chosen seller. Every consumer values each item at v , with a combined worth of w . Complements are incorporated by $w > 2v$, independent products are given by $w = 2v$ and substitutes are given by $w < 2v$.

Shelegia (2012) shows that if the market is controlled by a monopolist, the price of good a and b will satisfy $p_a \in [v, w - v]$ and $p_b = [w - p_a]$, if the products are complements. This yields a profit of w . Alternatively, if the products are substitutes, the price of both products is equal to $w - v$,

which generates a profit of $2(w - v)$. However, the more interesting case arises when the firm faces a competitor.

Consider first the case with complements. Firms follow a mixed strategy equilibrium that mirrors most of the features of a single product pricing equilibrium, except prices are now coordinated. There also exists a cut-off value between the highest price of one product in terms of the price of the other. This pricing condition elicits four possible variants, depending on the relationship between the price (p_i) and v (Shelegia, 2012). In the first case, labelled *strong complements*, both products are priced above their individual valuation and satisfy $p_a = w - p_b$ and $p_b = w - p_a$. In the second case, labelled *weak complements*, both products are priced below their individual valuation and must satisfy $\bar{p}_a = \bar{p}_b = w - v$. The remaining two cases, in which one product is priced above its individual valuation and the other product is priced below its valuation, cannot be solved. In addition, the fraction of *shoppers* determines whether prices follow weak or strong complements.

In the case of *strong complements*, which occurs when the fraction of *shoppers* is large, increasing the proportion of consumers who search both firms causes the upper bound of the price distribution to decline and prices become less dispersed. More importantly, the prices of the two products are negatively correlated. One product will be priced relatively cheap, whilst the other is relatively expensive. Shelegia (2012) also shows that firms can earn a higher payoff than in the one-product equivalent model but bundling the two products leads to lower profit. In particular, for a fixed proportion of *shoppers* (θ), both *shoppers* and *captives* are worse off in the two-product environment than in the single product case.

Under *weak complements*, which arises when the fraction of *shoppers* is small, Shelegia (2012) finds that firms continue to earn higher payoffs by selling the two products separately. However, relative to strong complements, the highest possible price will be higher and the lowest possible price will be lower, generating more price dispersion. The strength of the negative correlation between the product's prices is weaker in the case of weak complements because a high price for one product is not always accompanied by a low price for the other product.

In the case of independent products or weak substitutes, pricing is far less coordinated and resembles two separately priced goods. Each product is priced randomly without correlation between the products. In the case of perfectly independent items ($w = 2v$), the highest price in each distribution is the consumer's willingness to pay, v .

Zhou (2014) develops a model with endogenous sequential search in which two independent products become '*purchase complements*' for the consumer. The idea is that it is costly to visit multiple stores. Therefore, consumers prefer to buy all of their shopping at one location, which introduces co-pricing between unrelated products. Consider the following example. Firms sell the same two independent products, where independence implies that a price change on one product

does not affect demand for the other. The consumer visits one store and finds a very low price for product 1 but a high price for product 2. The consumer could choose to visit the rival for a better price on product 2 but the search cost outweighs the expected price saving. Therefore, both goods are purchased from the first store. If the price of product 1 was higher, the cost of visiting the second store would be distributed over the expected savings for both products and the consumer may find a second search worthwhile. Therefore, offering a discount on one item can cause consumers to buy other products from the store because of economies of search (Zhou, 2014).

In the model, two firms sell the same two independent and horizontally differentiated products to a unit mass of consumers, who hold unit demand for each item and buy both products but not necessarily from the same seller (Zhou, 2014). Buyers vary in their valuations of each product and a *match value* represents how well the product matches the consumer's ideal product. The valuation is drawn randomly for each product from the same joint distribution function, $F(u_1, u_2)$, where u_i denotes the realised match value of product i to this individual.

The game unfolds over two periods. In period one, sellers choose prices. In period two, buyers commence sequential search by choosing one of the firms at random. Visiting one firm incurs a search cost ($s \geq 0$) that is independent of the number of products examined at the store. The consumer may buy both, one or neither of the products at the first firm and can choose to visit the second firm, which incurs a second unit of the search cost.

The game is solved using Symmetric Perfect Bayesian Equilibrium. Consumers maximise their surplus, which is defined as $u_i - p_i$ for product i , where p_i denotes the price. Costless recall is assumed such that if a consumer visits a second firm, the prices and products at the first firm are always available to buy without incurring an additional search cost. As a result, consumers immediately buy both or neither of the products from the first firm.

Zhou (2014) shows that multi-product retailing creates new incentives for firms and induces a '*joint search effect*,' which states that a price reduction on one product can stimulate sales of a second product by deterring further search. To formalise this idea, consider the optimal shopping rule for a consumer seeking one product and suppose that the first firm's product has a match value x . Given that prices are symmetric across the firms,⁹ the only gain from further search is the anticipated increase in the match value:

$$\zeta_i(x) \equiv \int_x^{\bar{u}_i} (u_i - x) dF_i(u_i) \equiv \int_x^{\bar{u}_i} (1 - F_i(u_i)) du_i \quad (2.11)$$

Conditional on being offered products i and j with match values (u_1, u_2) , the consumer will only stop searching at the first firm if:

⁹Symmetric pricing arises from the assumptions required by the solution concept.

$$\zeta_1(u_1) + \zeta_2(u_2) \leq s \quad (2.12)$$

This produces an acceptance set of combinations (u_1, u_2) for which a second search is not undertaken. The frontier of this area defines the point of indifference between further search and stopping and can be expressed as $u_2 = \phi(u_1)$, where $\phi(\cdot)$ solves: $\zeta_1(u_1) + \zeta_2(\phi(u_1)) = s$. The frontier is also decreasing and convex. Therefore, if the utility from one product increases, the utility required from the second product to prevent further search, decreases. If one of the products (j) at the first store provides the highest possible utility level for the consumer, there is no benefit to further search for that product. The only possible gain from further search is to locate a better version of the second product and the indifference condition simplifies to $\zeta_i(u_i) = s$.

Reducing the price of one product has two effects. Firstly, the probability that consumers who search both firms buy from the firm, increases. Secondly, the discount reduces the probability that a consumer searches again, which stimulates sales of other products. This is the ‘joint search effect.’ The ‘joint search effect’ also survives when firms advertise their prices to consumers at the beginning of the game. In that case, prices are publicly observable and influence consumers’ search behaviour. However, the idea that reducing the price of one product can stimulate demand for other products, continues to pass through.

Zhou (2014) also demonstrates that firms can earn lower profit when search costs are higher because of two potentially competing effects. Firstly, *ceteris paribus*, the amount of search declines, which incentivises firms to raise prices. Secondly, the amount of consumers that lie on the reservation frontier also changes and the direction of this change depends on the assumed distribution in the model. In many cases, the reservation frontier grows, which incentivises firms to reduce their prices. For several distributions the incentive to cut prices dominates.

Petrikaite (2016) studies a different kind of multi-product retailing in which a monopolist sells two different varieties of a single product to a unit mass of consumers, who only want to buy one of the products. Consumers incur a cost to search for each product within the store, rather than incurring a search cost to visit the store. This allows Petrikaite (2016) to focus on the incentives for a firm to hide some of its products, to influence the order and intensity with which consumers search the store. In equilibrium, consumers consider products in ascending order of search cost. Therefore, in some senses this model bridges Wilson’s (2010) idea of directed search with the growing literature on multi-product retailing.

Products vary in quality for which consumers vary in their appetite. The value that consumer l places on product i is determined by a match value, ϵ_{il} . The utility from consuming a product for individual i is, therefore, the match value minus the price paid: $u_{il} = \epsilon_{il} - p_i$. The distribution of consumers’ match values is known to firms but the precise values for individual consumers are not

observable. Consumers search for products sequentially and the match value is an independent and identically distributed variable across the products and consumers, over the range: $[0, \bar{v}]$. The strategy for the monopolist is to devise a vector of prices and search costs for the two varieties: $\{p_1, p_2\}, \{s_1, s_2\}$ that maximises profit, where s_i denotes the search cost placed on product i .

Petrikaite (2016) demonstrates that the least obfuscated product, which consumers consider first, must have a higher price because the highest search cost acts as a screening device by inflating the price that buyers are willing to pay for the first product. Therefore, increasing the search cost on one of the products allows the firm to earn a higher payoff, which replicates some of mechanics in Wilson's (2010) duopoly framework.

It remains to consider which of the two differentiated products will be obfuscated. When the match value distributions are identical across products, neither product is a better candidate for obfuscation. However, when one product is preferred, in the sense that a higher frequency of consumers value it highly and a lower frequency of consumers place a low value on the product, a monopolist will charge a higher price for the preferred product and obfuscate the lower quality item. This stimulates sales of the higher quality good by reducing consumer's willingness to search for the discounted item. Therefore, it is profitable to offer a pair of high and low quality products but for a different reason to Ellison (2005).

The core dynamics of the model survive when prices are advertised to consumers at the beginning of the game but the match values remain unobservable. In addition, if neither product is preferred by consumers but one is more costly to produce, the costly product will be obfuscated. Obfuscation strategies are also unaffected by the introduction of a symmetric competitor, conditional on zero costs and zero recall between firms. In the duopoly case, visiting one store incurs an exogenous search cost s_f , which is the same across firms, and appraising each product incurs a store-chosen cost, s_p . The main insight is that both firms obfuscate one of their products, for the same reason as the monopolist in the main model. However, obfuscation may disappear entirely if prices are now advertised to consumers before searching.

Rhodes (2015) also studies multi-product retailing and allows firms to advertise the price of one of their products before consumers decide whether and where to search. Again, consumers' heterogeneous tastes are represented by a match value v_j for each product j . Consumers know their valuation of the products and wish to buy at most one unit of each but they do not know the prices until they incur the search cost.

In a monopoly environment without advertising, the seller chooses only the prices for each of their n products at the beginning of the game. In the second stage, consumers decide whether to visit the monopolist, which incurs a search cost and allows consumers to examine all of the prices. Therefore, similar to Zhou (2014), the search cost is best interpreted as a transport cost because in-store shopping is free. Rhodes (2015) demonstrates that the price a firm will charge is

positively related to the exogenous consumer search cost because when search costs are higher, some of the consumers rationally decide not to search the firm and consumers who are screened out by the high search cost have a low valuation for the product. Therefore, the average valuation of the products by consumers who search the firm is higher, which induces a higher price.

Prices will also be lower when the monopolist offers more products because a larger store will be able to attract more consumers, as their search cost is distributed over more products. However, similar to the case where the search cost declines, the average valuation of the products for consumers who search will now be lower and demand will be more price elastic. This incentivises the firm to lower the prices across the store to ensure that these consumers buy. The monopolist also earns a higher payoff when it offers more products, for two reasons. Firstly, for any given price, more consumers will choose to visit the store because the search cost is absorbed over more products. Secondly, a firm with more products can lower their prices to induce more consumers to visit the store and buy. Therefore, larger stores with more products should charge lower prices than their smaller counterparts because the demand they face is more price elastic.

The idea that a change in search costs or the size of the store will cause the amount of consumers that visit a firm to change, and critically cause the average elasticity of demand within the searching population to change, is reminiscent of Ellison's (2005) model of upgrade pricing. However, the dynamics are very different. In Ellison (2005), many firms supply a low and high quality product but competition is suppressed for low quality items because a price cut causes the most price sensitive buyers to switch to the firm, and these consumers will never upgrade to the higher quality item. In Rhodes (2015), prices influence both the amount of consumers that search it and the average valuation of the products from consumers who search.

These effects become more pronounced when the firm can advertise one of the prices to consumers before the game begins, at cost $k > 0$. The timing is partly sequential as the firm commits to an advertising strategy in period zero and must choose the price of any advertised product. In period one, the firm chooses the prices for unadvertised products before consumers decide whether to search and make their purchasing decisions. If the firm advertises one price, consumers form rational expectations of the prices of other products at the store. If the firm does not advertise, consumers read into their decision.

The prices of advertised and unadvertised products will be positively correlated and consumers correctly infer that 'hidden prices are higher prices,' which connects to the add-on framework of Gabaix and Laibson (2006). The intuition for a positive correlation is similar to the no-advertising game. When a firm sets a higher price for an advertised product, consumers will rationally infer that other products at the store will be expensive. Therefore, consumers with low valuations drop out of the market and the firm can charge a higher price. The price of the advertised product is also lower than the monopoly price for three reasons. Firstly, the presence of

search costs heightens consumers' sensitivity to prices. Secondly, when the advertised product is cheaper, more consumers find it worthwhile to visit the firm, which can boost sales of other goods. Thirdly, a low advertised price causes consumers to expect that other items in the store will be competitively priced, introducing a reputation effect for the firm.

The positive correlation between advertised and unadvertised prices differs markedly from the add-on pricing models discussed in Section 2.2.1. In that literature, firms sell one product at a discounted price to stimulate sales of high margin add-ons. However, in Rhodes' (2015) search model, prices are positively correlated because the price a firm charges determines the price sensitivity of its clients. The framework is also closely aligned with the '*joint search effect*' described by Zhou (2014). However, in that story a price cut on one product causes more consumers to cease searching and buy other products in the store. In Rhodes' (2015), a price cut on one product causes more consumers to visit the monopolist, which boosts sales of other items. Therefore, in Rhodes (2015) the amount of consumers who search the firm changes, whereas in Zhou (2014) the amount of consumers that terminate their search at the firm changes.

More generally, Rhodes (2015) finds that if the monopolist decides to advertise, prices across the store are lower. Therefore, the question of whether a firm will advertise is determined by the trade-off between the additional revenue that advertising generates and its cost. Rhodes (2015) also shows that the monopolist's gains from advertising are non-monotonic in the amount of products it carries. Advertising is beneficial as the initial number of products increases but not profitable for larger stores with many products because the proportion of consumers who choose to search is already higher. Larger stores are also more likely to create a reputation for low prices, whereas firms with fewer products advertise a discounted product and charge high prices for non-advertised products.

In the second part of the study, Rhodes (2015) explores a duopoly version of the model. Two firms sell the same two independent products to consumers, who wish to buy one of each item. Consumers know their valuations of the products but they do not initially know the prices. Similar to Varian (1980), some of the consumers are *shoppers* who have no preference for either firm but the remainder are *loyal* and only visit a randomly chosen seller.

At stage zero, firms choose whether to advertise one product and its price. At stage one, firms choose unadvertised prices without knowing the stage zero decisions of their rival. At stage two, consumers learn any advertised prices and make their search decisions. Loyal consumers only visit one seller but non-loyal consumers may search both firms, with a search cost $s > 0$ for each firm visited. Non-loyal consumers can also return to a firm for a fee, $r \in (0, s)$.

If firms do not advertise, non-loyal consumers randomly choose which firm to search first. If both firms advertise the same product, non-loyal consumers visit the firm with the cheapest advertised price. If firms advertise different products, non-loyal consumers buy the advertised

product of each firm. The prices that a firm sets under advertising are random draws from a distribution due to the trade-off between offering a low price to attract non-loyal consumers and charging a high price to exploit loyal consumers (Narasimhan, 1988; Varian, 1980). The decision to advertise is also randomised. Each firm chooses to advertise with probability $1 - \alpha$, where α is a decreasing function of the cost of advertising (Rhodes, 2015).

In an extension, Rhodes (2015) shows that the results survive when a consumer may buy multiple units of the products. In addition, when the two products are heterogeneous in terms of their interest for consumers, in the sense that consumers have higher valuations for products that they are most interested in, a monopolist will advertise the price that is of most interest to consumers. This connects to Petrikaite's (2016) framework in which less preferred products are obfuscated.

2.2.5 Frequently Changing Prices

Models of price competition predict price dispersion and frequently changing prices as outcomes of a competitive market when some consumers do not buy at the lowest price (Conlisk et al. (1984); Narasimhan (1988); Sobel (1984); Varian (1980)). Selecting a price deterministically exposes the firm to a rival that can marginally undercut their price and steal all of the informed consumers. In practice, however, the assumption that all informed consumers will instantly switch to a provider with a marginally lower price is too strong, which casts doubt on the need for firms to change their prices frequently. One alternative explanation for firms to update their prices is that more consumers learn the prices available over time. Therefore, a firm can adjust its prices to prevent too many consumers from becoming sophisticated shoppers. This approach to obfuscation differs significantly from the preceding discussion where consumers do not change over time.

Carlin and Manso (2011) explore a monopolistic market in which consumers become increasingly sophisticated over time. We can divide the consumers in their model into three types. Denote the proportion of consumers that understand all of the prices in the market at time t as x_t . Some consumers (x_0) will be informed and aware of all prices at the beginning of the game. The remainder ($y_0 = 1 - x_0$) are initially uninformed but over time some of these individuals become informed. That is, $\frac{\partial x_t}{\partial t} (> 0) = -\frac{\partial y_t}{\partial t}$.

Carlin & Manso (2011) show that this creates an incentive for the seller to systematically refresh prices, at cost, to soften competition. The frequency with which firms revise their promotions, however, depends on both the cost of changing the menu and the rate of consumer learning. The model abstracts from price setting and the expected payoff for the monopolist is written:

$$\pi(x_t, y_t) = ax_t + by_t \quad (2.13)$$

a and b represent the exogenously determined payoffs that a firm receives from informed and uninformed consumers respectively, where $a < b$ such that informed consumers are less profitable for the firm. The optimal strategy defines a vector of optimal time points, $T = (t_1, t_2, t_3, \dots)$, at which the firm should refresh its prices.

Carlin & Manso (2011) consider two types of learning processes: Independent and interactive learning. Independent learning refers to uninformed individuals learning independently and interactive learning refers to situations in which consumers learn as a community from each other. Under both specifications, the optimal time between refreshes is increasing in the cost of updating prices and decreasing in the profit difference that a firm earns from the two types of consumers. However, the relationship between the refresh rate and the rate of consumer learning is non-monotonic. When consumers learn quickly, the gains from refreshing the menu are short-lived. Therefore, frequent refreshes are costly and generate short term gains (Carlin and Manso, 2011). When consumers learn slowly, the gains from refreshing the menu last longer but the firm need not refresh the menu as often. The highest refresh rate occurs at an intermediate rate of consumer learning and all price changes are socially wasteful.

The learning mechanism also determines the relationship between the proportion of initially informed consumers and the refresh rate. When consumers learn independently, firms increase the duration between refreshes when more consumers are initially informed because the additional rent that a firm can achieve from refreshing their menu is smaller. Therefore, improving consumer sophistication will have an additional benefit by softening the obfuscation incentives of firms. If individuals learn interactively, the association between the initial proportion of informed agents and the refresh rate is non-monotonic. If the initial proportion of informed individuals is very high, the gains from obfuscation are small. Conversely, if the initial proportion of informed individuals is very small, learning will take a long time because individuals only become informed when they come into contact with an informed individual.

Introducing competition into the model also causes firms to reduce their obfuscation because the market separates for each type of buyer. Consistent with Bertrand competition, each firm charges informed buyers a price of zero and charges uninformed buyers their reservation price. However, the guaranteed market share of a firm approaches zero as the market becomes perfectly competitive, which yields zero rent for the firm and eradicates price changes. One critical assumption in this analysis is that consumers who cannot understand prices will buy randomly, which I revisit in Section 2.4 and throughout this thesis.

2.3 Evidence

Despite the prevalence of spurious price complexity in retail markets, the evidence of obfuscation behaviour in practice, is limited. One explanation is that many obfuscation strategies can be recast as by-products of competition. Shrouding some attributes of a product may be justified as an attempt to provide consumers with only the most important information or arise from high advertising costs. Tariff proliferation may arise from the existence of heterogeneous consumer preferences. Prices may also fluctuate as the outcome of an entirely competitive ‘mixing equilibrium.’ Therefore, it can be difficult to distinguish between anti-competitive obfuscation and the consequences of competitive market interactions.

In Section 2.3.1, I survey the existing empirical literature. The experimental literature is considered in Section 2.3.2. I do not include the growing literature on consumer search patterns online, which has gained traction due to the availability of detailed consumer data. That branch of studies is dominated by work within the marketing discipline and is a suitable candidate for a future inter-disciplinary review. I also do not include the established literature on false or misleading advertising, which is an intimately related but distinct literature from pricing and strategic obfuscation.

2.3.1 Empirical Evidence

Several theoretical models predict that firms can profitably increase the number of dimensions on a contract or the quantity of products they offer, to make consumers’ decisions more taxing. However, in some markets, particularly during the 1980’s and 1990’s when technology was undergoing significant development, firms offered very few variations of their products. Miravete (2006) reports evidence from the U.S. telecoms market to develop the argument that firms do not need to obfuscate their products in the absence of competition. In the 1980’s and 1990’s these firms operated as regional monopolies. Therefore, firms are able to extract nearly all of the willingness to pay of each consumer using a non-linear tariff, with a very limited role for obfuscation. Miravete (2006) also contends that the consumers who engaged with this market during its infancy, exhibit a low willingness to pay for telecoms services. Therefore, the optimal pricing strategy involves excluding these ‘*low types*’ altogether and maximising the informational rents that can be extracted from individuals with a high willingness to pay.

The incentive to obfuscate in this market is limited because the firm does not face a competitor. Therefore, the monopolist need not consider the additional gains of confusing consumers in terms of preventing consumer switching. The only trade-off exists between using more tariffs to increase the rent from each consumer and using fewer tariffs to curtail operating costs. However, these incentives differ markedly in markets with several firms, which is the focus of Ellison and Ellison

(2009) and Miravete (2013). Ellison and Ellison (2009) use data from a price comparison website to analyse obfuscation and pricing strategies of computing memory sellers and Miravete (2013) explores the emergence of competition in the same 1980's market for telecoms carriers.

Ellison and Ellison (2009) focus on two main obfuscation strategies. Firstly, loss leader pricing: Do firms advertise a low quality and low price product, before encouraging consumers to switch to a higher quality and more expensive product at the point of sale? This follows the 'quasi-bait-and-switch strategy' outlined in Ellison (2005). Secondly, do firms deliberately inflate search costs? They examine an oligopoly of computing memory suppliers, who list on a price search engine: PriceWatch.co.uk. The products are minimally differentiated and there are many small sellers. The price data is also rich and products can be easily categorised in terms of quality. The search engine is also the primary channel of sales for the firms, prices are updated frequently and each seller must update their price manually. This prevents firms from manipulating the search engine, which occurs frequently when prices are drawn from a supplier's website. Consumers are also particularly price sensitive. Ellison & Ellison (2009) estimate that a 1% increase in price can induce a 20% reduction in sales. Retailers on the website are listed in ascending price order, which should enable consumers to quickly buy at the lowest price. However, this is not the case.

The data consists of a dynamic annual series of hourly price and sales statistics between 2000-2001. Products are categorised into three quality groups to provide a classification system for the 'bait' and 'switch' products. They also obtain firm-level sales and cost data. A General Method of Moments (GMM) framework is applied to compute the elasticities of consumer demand for the respective products of each firm, combined with instrumental variables to prevent endogeneity biases amongst the prices of firms.¹⁰

Their results provide strong support for the two channels of obfuscation. Firstly, firms increase consumers' search costs by concealing product and price information, and bundling extremely low product prices with high shipping charges. Pricewatch has previously attempted to address such practices by introducing a maximum price for standard shipping, encouraging suppliers to display delivery prices alongside product prices and warning consumers that firms who hide their prices may be more expensive. PriceWatch also started to list sellers in ascending order of their 'total prices,' inclusive of shipping and handling fees. This generated some improvements in transparency but only for standard shipping methods. Faster shipping options, with high prices, remain hidden.

In a somewhat Spiegler (2006) fashion, sellers also bundle low total prices with unfavourable terms and conditions, including reduced warranties and high return fees. Firms also include add-

¹⁰GMM would otherwise fail to account for the anticipated positive correlation between the price series' of each firm, where a change in the price of one firm is likely to cause other firms to change their prices in the same direction. A further endogeneity issue is the negative correlation between the price of a firm's product and their sales (Ellison and Ellison, 2009).

ons on every order by default and discriminate across the market by offering competitive prices on the search engine that are hidden on their own website.

For the second obfuscation channel using quasi-bait-and-switch behaviour, Ellison & Ellison (2009) find that firms offer damaged goods at lower prices, before encouraging consumers to upgrade to superior products. In particular, sellers highlight the inferior quality of the attention grabbing products once the consumer has committed to purchasing from the firm. The price of low quality modules also has a significant impact on demand. A firm that falls from first to seventh in the price order for low quality memory can expect an 83% reduction in sales of low quality items. Prices of low quality products also vary more frequently than other qualities and act as loss leaders as described by Ellison (2005): Sales of higher quality modules increase significantly when the firm has a more competitive price in the listings of low quality goods.

One novelty in Ellison's (2005) theoretical model is an adverse selection effect that limits price competition on low quality modules. The intuition follows that a firm who reduces the price of their low quality product will attract more buyers, but a larger share of these buyers will not upgrade to the profitable higher quality item. Ellison and Ellison (2009) indicate support for this mechanism: Proportionally fewer customers of the lowest priced firm choose to upgrade. Further analysis of the price-cost data also indicates that retailers sustain significant mark-ups in the range of 16% and 27% for medium and high quality units, respectively. Most concerning, the majority of consumers fail to buy from the cheapest supplier.

In a second study of the U.S. market for telecoms carriers, Miravete (2013) uses dominated tariffs to measure the *fogginess* of pricing schemes and studies the impact of a transition from monopoly to duopoly between 1984-1988. Due to licencing issues, local carriers enjoyed temporary monopoly status. However, a second firm was always scheduled to enter, once approved by local legal authorities. Therefore, the timing of entry into a local market is independent of both pricing by the incumbent and the timing of entry in other regions (Miravete, 2013).

Tariffs during this period are relatively straight-forward and most contracts consist of a fixed access charge, inclusive minutes and a fixed price for each unit of additional usage. Prices also vary across peak and off peak periods. This differs from the mature online market studied by Ellison & Ellison (2009). The paper also differs in focus from Ellison & Ellison (2009). They focus on the obfuscation strategies employed by oligopolists when a search engine attempts to increase transparency but Miravete (2013) focuses on the use of deceptive pricing tariffs when firms lose monopoly power.

Three measures of *foggy pricing* are proposed: The quantity of dominated tariffs offered by each firm, the proportion of tariffs that are dominated and the complexity of the non-dominated tariffs. Miravete (2013) defines dominated tariffs as contracts that are never cheapest for the consumer,

irrespective of consumption. That is, it is possible to find a contract, or a combination of contracts, that is always cheaper.

Two datasets provide detailed tariff information of the carriers in the largest 100 U.S. cities. This includes information on tariffs during the period of structural change and further information in 1992. This facilitates analysis of the immediate and medium-term implications of introducing competition on firms propensity to obfuscate. Actual consumption statistics, however, are not available and consumption profiles are simulated to identify dominated tariffs.

Initial results indicate a multi-directional impact of competition on the use of *foggy pricing*. During the monopoly phase, 60% of firms offer no more than three tariffs, with one third offering only one. However, in the post-entry period, 62% of existing suppliers and 54% of new entrants offer 3 or more tariffs and very few firms continue to offer only a single tariff (3% of incumbents and 9% of entrants). Miravete (2013) shows that existing firms tend to increase the number of *foggy* tariffs when a second firm enters. Before entry, 22% of sellers offer only one *foggy* contract. Post entry, 27% of sellers offer only one *foggy* contract. More concerning, before entry, 20% of incumbents offer two *foggy* options, which increases to 35% when a competitor arrives.

Competition also causes firms with few dominated options to reduce the proportion of dominated tariffs but firms initially offering a high proportion of dominated contracts increase the proportion of dominated tariffs. More specifically, when firms offer just two tariffs in the monopoly period, 13.11% of tariffs are dominated. This figure declines with competition to 3.84% for the existing firm and 6.13% for the new firm (Miravete, 2013). However, the fraction of firms that offer three dominated tariffs out of five, increases from 1% to 8.07% and 2.77%, for incumbents and entrants, respectively. Therefore, firms respond to competition by offering more dominated tariffs but the proportional share of tariffs that are dominated is ambiguous.

Incumbents and entrants also follow different strategies. Incumbents are more likely to increase the *fogginess* of their prices as a defence mechanism against competition but new entrants offer more transparent prices.¹¹ The conclusions of Miravete (2013) cast doubt on the incentives for firms to respond to competition by increasing the *fogginess* of their pricing. However, several remarks are required. Firstly, the empirical analysis considers historic telecoms markets, which operate very differently from modern retail markets. Secondly, modern retail markets typically comprise several competitors, rather than two. Thirdly, it is also not clear that the pricing strategies adopted by incumbents are completely independent of competition if the monopolists' created their price tariffs whilst anticipating entry from the second carrier (Miravete, 2013).

With the advent of internet technologies, one might expect that consumers would be better placed to avoid choosing dominated contracts in modern markets. However, recent evidence from Wilson and Waddams-Price (2010) indicates that this need not be the case. Wilson and

¹¹These results are robust to uncertainty regarding future consumption for consumers.

Waddams-Price (2010) analyse consumer switching decisions in the domestic electricity market across England, Scotland & Wales. Historically, consumers were geographically tied to the service of their regional provider, with 14 regions in total. However, in 1999 the market was liberalised. This allows consumers to choose their supplier and firms to operate in several regions. Within 10 years approximately half of consumers switched supplier. In the short term, market concentration decreased significantly but many suppliers could not survive in the long run (Wilson and Waddams-Price, 2010).

Electricity is suited to economic analysis because electricity is the same regardless of the supplier. Their data is drawn from two surveys. The EA survey records responses of 1427 consumers during the year 2000. Only 523 consumers had already switched their supplier due to the short time period between market liberalisation and the survey. Approximately 75% of consumers who had switched provided detailed information regarding their decision. In the Centre for Competition Policy survey, data for 2027 consumers is recorded. 370 individuals had switched supplier and 250 provided the full information of their switch and the reason for switching. This enables Wilson & Waddams Price (2010) to identify consumers who switched purely for price savings. Consumers who reported non-price explanations for switching were excluded.

At the heart of their analysis is a comparison between the actual gain from switching obtained by the consumer and the maximum gain available at the time of switching.¹² The theory underpinning their model is as follows. Suppose that consumer i switches from their original supplier (o) to a new supplier (n). The change in consumer surplus (CS) can be written:

$$\Delta CS_{ion} = CS_i^n - CS_i^o \approx [u_i(C_i^n) - E(C_i^n; T^n)] - [u_i(C_i^o) - E(C_i^o; T^o)] \quad (2.14)$$

T^n and T^o denote the tariffs offered by the new and original suppliers, respectively. The utility from consuming from firm o is the difference between the utility from consumption and the price of that consumption. The difference between the two utility levels indicates the change in consumer surplus. The maximum gain from switching is x_{ion}^{sw} . This allows the cost of a fixed consumption level to be compared across the old and new tariffs:

$$\Delta CS_{ion} \leq x_{ion}^{sw} = E(C_i^n; T^o) - E(C_i^n; T^n) \quad (2.15)$$

The highest possible increase in consumer surplus from switching is also visible based on the consumption plan and the range of alternative tariffs in the market. Define this value as: x_{io}^{max} . Similarly, the gains for a consumer that switches to a randomly chosen supplier is given by: x_{io}^{mean} . These values are written explicitly as:

$$x_{io}^{max} = E(c_i^n; T^o) - \text{Min}_{k \in S_i}, E(c_i^n; T^k) \geq 0 \quad (2.16)$$

¹²Future consumption uncertainty can be excluded because consumers switch conditional on a specified level of expected usage.

$$x_{io}^{mean} = E(C_i^m; T^o) - Mean_{k \in \{S_i \setminus o\}}, E(C_i^m; T^k) \quad (2.17)$$

S_i defines the range of alternatives in the market for consumer i . If consumers switch to the cheapest tariff, $x_{ion}^{sw} = x_i^{max}$.

Wilson & Waddams Price (2010) find that very few consumers (8% - 20%) switch to the cheapest supplier, which translates into an average saving of £17 – 23 per year, rather than the £41 – 74 saving available. More concerning, 17 - 32% of switchers move to a more expensive tariff, which generates an annual loss of £14 – 27. Between 19% and 48% switch to a dominated tariff. The presence of an agent in the switching process, such as a cold caller or firm representative, has no impact on savings but reduces the probability that the consumer switches to the cheapest tariff. The optimality of consumer switching also appears to be independent of the number of competitors.

Frank and Lamiraud (2009) also investigate switching behaviour across health insurance providers in Switzerland. Consumer participation is guaranteed because citizens must purchase at least basic cover. Individuals may buy additional cover, although not necessarily from the same supplier.¹³ The basic policy is also standardised across firms and local authorities distribute annual reports of prices to consumers. Between 1994-2004, the national market became more concentrated but local markets (Canton's), of which there are 26, witnessed an increase in consumer choice. It is therefore surprising that prices in this market remain stubbornly high with an estimated dispersion of 80% of the average contract price for basic cover. Switching is also below 6% annually (Frank and Lamiraud, 2009).

Their main focus is to distinguish if the low rate of switching and sustained price dispersion can be supported by theories of: Expected utility maximisation with costly search, consumer choice overload or status quo biases. In addition, they study the effects of more providers operating in a Canton. Their data combines several sources. Insurance market information between 1997-2000 is obtained, which includes information on price plans, the insurer, the time period and the Canton. This permits an estimation of price dispersion. A survey undertaken by the Swiss Federal Office for Social Insurance is also included. This contains information on 2512 households, including their current plan, reasons for choosing the plan, switching behaviour and propensity to switch in the future, which is compiled into a panel dataset.

The specific survey information from households is then connected to information from the insurance market. This links individual's contracts with prices. Over 70% of survey respondents had access to more than 50 insurance plans in 2000. 15.2% reported that they have switched within the preceding four years but only 6 households changed more than once. 20% stated an intention to switch in the future (Frank and Lamiraud, 2009).

¹³Frank & Lamiraud (2009) do not explore the co-pricing strategies between basic and supplementary insurance products. This is one direction for further research in this market.

At the heart of their model is an econometric estimate of the propensity for a household to switch, y_{nt}^* , which is given by:

$$y_{nt}^* = x'_{nt}\beta + (dp)_{nt}\eta + O_{nt}\gamma + S_{nt}\lambda + year_t\alpha + c_n\rho + \epsilon_{nt} \quad (2.18)$$

Switching occurs whenever $y_{nt}^* \geq 0$, which leads to $y_{nt} = 1$ as a binary variable to indicate switching. x'_{nt} captures the information on each household, including time varying parameters. $(dp)_{nt}$ is the estimated potential savings available from switching, given by an adjusted standard deviation of prices for the Canton. O_{nt} indicates the number of available insurers in the Canton. S_{nt} is a binary variable that takes the value 0 if the individual opts only for basic cover. This separates individuals with basic cover from individuals with additional protection. Time fixed effects are included through $year_t$.

They find that switching and intention to switch are markedly lower in Canton's that contain more insurance options. Therefore, when consumers face more choices, engagement declines. Consumers who did switch save an average of 15.9% but the likelihood of a consumer switching decreases with the length of time they spend on their current contract. Households that seek advice from intermediaries save approximately 10%. Therefore, it is not an absence of price savings underpinning the low switching rates.

Survey information reveals that many households choose the same plan as family members and friends, with 25% of individuals stating that their objective is not to obtain the cheapest contract. Frank and Lamiraud (2009) also find that many households who are not satisfied with their current insurance, do not intend to switch. In Canton's where there are fewer than 50 available options, 33.7% reported an intention to switch. This figure declines to 21.9% in Cantons' with more than 50 alternatives. Therefore, as choice increases, consumers become increasingly attached to their current contract, even when dissatisfied.

In a more recent study, Muir et al. (2013) analyse a unique dataset that contains unusually rich information on both the pricing behaviour of firms and the purchasing behaviour of consumers, in the Portuguese market for driver training courses. In Portugal, learner drivers purchase a standard driving package from an approved school. This covers lessons and a first attempt at the theory and practical driving examinations. If the student passes, no further fees are due. However, if the student fails any component, which nearly half of students do, they must take a retest for which the driving school charges additional fees (Muir et al., 2013).

They study the pricing behaviour of the school's and specifically the co-pricing strategies between basic course packages and retest fees. The market fits closely with the add-on product environment of Gabaix and Laibson (2006). There is a clear base product and costly add-ons, which some consumers do not foresee. The base course is homogeneous and students are tied into their first driving school for the entire course and retests. This enables firms to profit from

consumers that neglect part of the fee structure. However, in contrast to Gabaix & Laibson (2006) where add-ons are unavoidable, retests may not be required. This permits over-confidence as a rational explanation for learners to ignore retest fees.¹⁴

Muir et al. (2013) obtained 782 student surveys, completed at the end of their theory examination (but before they learned the outcome). Students predicted their probability of passing the theory component with a high degree of accuracy but significantly overestimate their chances of passing the practical. This compounds with the traditional explanation for consumers to neglect additional fees, based on limited foresight. Surprisingly, only 78.5% of students were certain that retest fees existed and conditional on being aware of the existence of retest fees, only 66.5% of students knew the price.

Price data from 1119 schools on mainland Portugal is hand-collected by mystery shoppers. This covers 57,329 drivers across 420 schools in 158 municipalities, which acts as the geographic market for the purpose of market definition. Approximately 30% of the data is excluded due to truncation of students that had already begun or failed to complete the entire test process within the time frame.

Prices in this market, however, tell only part of the story and the main interest is the profit margin. To calculate this, Muir et al. (2013) estimate marginal costs using data from both driving schools and an industry authority. Descriptive analysis of the data reveals that retest fees represent a significant revenue stream for the industry and generate more than 20% of profit. Mark-ups on retests, measured using the Lerner index, are 86.5% and 56.6% for the theory and practical components, respectively. This is significantly above the average 27.7% mark-up on the basic course fee and for six schools the basic package is a loss making activity.

Muir et al. (2013) also find that the number of sellers has no bearing on retest mark-ups but exerts significant downward pressure on the prices of the basic course package. Therefore, whilst competition may benefit consumers who pass first time, students who require retests may not benefit from additional entry. Overall, this comprehensive study demonstrates considerable support for consumers neglecting add-on costs and provides robust evidence of loss leader pricing in the spirit of Gabaix and Laibson (2006).

The fundamental mechanisms at play are also different from those studied by Ellison and Ellison (2009). In their computer memory industry, low quality products draw consumers to the website, at which point consumers are encouraged to upgrade to higher quality modules. In the driving industry, low prices for basic packages are cross-subsidised by higher retest prices. Therefore, whilst all learner drivers consume the loss leader when learning to drive, firms dissuade consumers from buying the basic product in the market for computer memory.

¹⁴See Grubb (2015b) for a review of over-confidence in consumer choice.

2.3.2 Experiments

Field and laboratory experiments provide an alternative approach to analysing pricing and obfuscation strategies that overcomes the issue of incomplete and noisy empirical data. When designed correctly, experiments can isolate the precise forces at play in consumers' decision-making processes and test the impact of specific obfuscation and pricing strategies (Sitzia and Zizzo, 2011).

Morgan et al. (2006) provide some of the earliest experimental evidence that reducing the fraction of informed buyers will generate higher prices. In the laboratory, demand is automated by computers, which specifies the fraction of informed and uninformed buyers. Informed buyers are programmed to buy at the lowest price and uninformed buyers buy randomly. Participants play the role of firms and are randomly assigned to a market with 1 or 3 other competitors.

Their game unfolds over 90 periods. In the first and final 30 periods, 3 out of 6 of the consumers are informed. In the second 30 periods, 5 out of 6 are informed. At the beginning of each period, participants choose a price, conditional on the number of competitors and the fraction of informed buyers. They show that prices are higher when proportionally fewer consumers are informed. Average prices also increase with the number of competitors but the lowest price decreases. Therefore, informed consumers are better off whilst uninformed consumers are worse off. Price dispersion also increases with the fraction of informed buyers and the number of competitors.

In a related field experiment, Hossain and Morgan (2006) investigate price co-ordination between base goods and add-ons on an online auction platform. They auction 40 CD's and 40 Xbox games on eBay, where the listing includes an initial price for the product and a fixed delivery charge, both chosen by each seller. They analyse how variations in the distribution of the initial total price between the opening bid and delivery charge influences demand. Four price combinations (measured in U.S. Dollars) are considered, as summarised in Figure 2.1.

Treatment	Opening Bid	Shipping Fee	Total Price	Price Level Type
A	4	Free	4	Low
B	0.01	3.99	4	Low
C	6	2	8	High
D	2	6	8	High

Figure 2.1: Hossain & Morgan (2006)

Treatments A and B and treatments C and D have the same total opening price but the distribution varies. In A and C, the opening bid is the largest component. In B and D, the shipping fee is the largest component. Comparing bidder's behaviour in A with B, and C with D, reveals the impact of changing the proportional value of the shipping fee. The high type treatments (C & D) are also

\$2 more expensive on every dimension than the low type treatments (A & B). Comparing between treatment types captures the impact of changes in the total initial price.

For CD's, one product was unsold at the end of the auction under the low opening price type and five products went unsold under the high opening price. Therefore, offering a low opening bid increases the probability of a sale. This does not occur for Xbox games because all products sell. The explanation is that retail prices are usually much higher for Xbox games higher than CD's. For products that do sell, a higher total initial price generates higher final prices. For CD's, the average final prices are \$9.30 and \$12.21 for the low and high treatments, respectively. For Xbox games, the average prices are \$37.47 and \$39.01.

Under the low type treatments, revenues are also higher when shipping fees are a larger share of the total initial price. In the case of CD's, average revenue is \$7.54 and \$10.14 for treatments A and B, respectively. Removing the CD's that failed to sell, the price under treatment B is more than 20% higher than under treatment A. Therefore, inflating shipping fees and reducing the minimum bid price leads to higher revenue.

In the high type treatments, more CD's went unsold. Including products that did not sell, revenue for CD's did not vary significantly when shipping fees are a larger or smaller proportion of the total initial price, and average revenue is marginally higher in treatment C (\$11.63) than treatment D (\$11.26). However, when the unsold products are removed, treatment D generates higher revenue. Therefore, conditional on being sold, inflating the shipping fee and reducing the initial bid, increases revenue.

For Xbox games, all products sell and it is unsurprising that revenue is more than 10% higher under treatment D than treatment C. Therefore, for products where the initial total cost is high, relative to the value of the product, expected revenue is not markedly different between the two pricing strategies because of the reduction in sales. However, when the initial cost represents a significant discount, sellers should place more weight on shipping fees.

Hossain & Morgan (2006) also find that for CD's: Treatment B has an average of 0.6 more bidders than treatment A, and treatment D generates an average of 0.2 more bidders than treatment C. For video games: Treatment B attracts one additional bidder than treatment A, and treatment D attracts an additional 1.6 bidders than treatment C but only some of these differences are statistically significant.

The average quantity of bidders is also lower for CD's when the initial total price is higher. That is, fewer individuals bid under C (D) than A (B). This is consistent with the idea that individuals with a reservation price between the two total opening prices are screened out on CD's. This does not appear for Xbox games because few consumers hold a reservation price between these two levels, as Xbox games are more expensive (Hossain and Morgan, 2006). Their final result is that auctions that achieve more attention from consumers, in the form of more bidders, lead to a

higher end price and auctions with a low initial total price attract bids earlier than auctions with a high initial price.

Brown et al. (2010) address a complementary question: Is it profitable for sellers to disclose or shroud delivery charges? This directly tests the theoretical literature reviewed in Section 2.2.1. They borrow the account of an established seller of Apple iPod's on two online auction platforms: In Taiwan during 2006 and in Ireland during 2008. Sellers choose the price and the delivery fee but are ranked by the product price. To learn the shipping fees, consumers must manually scroll through the product page. Therefore, disclosed fees are given by including the charge in the title of the product listing and shrouding involves burying the information in the product listing.

In Taiwan, six iPod models that vary in their memory and colour were sold and six treatments were applied to each iPod. The treatments are summarised in Figure 2.2. All figures are given in Taiwan Dollars (TWD) with TWD 0.33 = USD 1 and OP defines the opening bid price.

<i>Taiwan</i>	OP: 750	OP: 750	OP: 600
	Low Shipping: 30	High Shipping: 180	High Shipping: 180
Disclose (D)	DL	DH	DR
Shroud (S)	SL	SH	SR
<i>Initial Price</i>	780	930	780

Figure 2.2: Treatment Types: Taiwan, Brown et al (2010).

Comparing Dx with Sx reveals the impact of shrouding effects. Comparing xL with xR reveals the impact of increasing the proportion of the total opening bid price that is placed on the shipping fee, whilst maintaining the same overall initial price. Comparing xR with xH highlights the effects of increasing the initial price, without changing the delivery fee. Typical delivery fees range from TWD 50 to TWD 250. Therefore, the low shipping fee (TWD 30) is very cheap and the high shipping fee (TWD 180) is reasonably high.

Movements in the total initial bid price were found to have no bearing on the outcomes of the auctions on the Taiwan platform in 2006. Therefore, treatment xR is omitted in the auctions in Ireland in 2008. The remaining treatments are summarised in Figure 2.3.

<i>Ireland</i>	OP: 0.01	OP: 0.01
	Low Shipping: 11	High Shipping: 14
Disclose (D)	DL	DH
Shroud (S)	SL	SH
<i>Initial Price</i>	11.01	14.01

Figure 2.3: Treatment Types: Ireland, Brown et al (2010).

Prices are denominated in Euros (€), with both opening bids for the main product priced at €0.01. Two auctions were listed each week for each of the four treatments, and five weeks of data were acquired, with 40 iPods sold. Again, delivery fees are within the range of the market at the 25th and 75th percentile (Brown et al, 2010).

The results of the two experiments are pooled for analysis. They find that the impact of shrouding delivery fees on revenue depends on the size of the delivery fee. When the fee is small, obfuscating reduces revenue. However, when fees are high, different outcomes prevail in different markets and they fail to reject the hypothesis of no significant effect of disclosure. The influence of changing the size of a disclosed shipping fee on revenue also varies across the markets with no statistically significant effect. Under shrouded fees, however, the situation is different and increasing the shipping price increases the final price by 5% in Taiwan and 7% in Ireland. Therefore, retailers who hide low delivery prices could earn higher revenue by increasing their delivery fee. They also find no evidence that changing the delivery fee influences the number of bidders.

The second part of their analysis involves developing a theoretical framework that encompasses their empirical observations. Similar to Gabaix and Laibson (2006), they have three types of consumers, where some pay attention to the total price but others do not. However, in contrast with Gabaix & Laibson (2006), all consumers know that delivery fees exist and consumers vary only in their expectation of the fee. Some inattentive consumers anticipate low delivery fees (*naive buyers*), others anticipate high delivery fees (*suspicious buyers*). Therefore, disclosing the fee causes some inattentive buyers to update the price that they are willing to bid.

Naive consumers, who expected a low price, will reduce their bids. However, suspicious bidders, who expected a higher price, will increase their bids. Disclosing the delivery fee is, therefore, revenue enhancing when the shipping charge is low because the increase in the bid values of suspicious bidders dominates the decrease in the maximum value that naive bidders will offer. When the delivery fee is large, disclosing the fee causes a larger reduction in offers from naive bidders than the increase in offers from suspicious bidders.

Brown et al. (2010) also report a natural experiment based on changes in the organisation of product listings on eBay in the U.S. Before 2004, delivery charges were only visible to consumers who searched through the individual product listings, which mirrors the shrouding strategy in the laboratory experiments. However, in 2004, eBay changed the search features to allow buyers to order products by the total price. Therefore, shrouded auctions refer to listings before 2004. Disclosed auctions refer to listings post-2004. They analyse the final selling prices of different types of silver and gold coins. Before the regime switch, increasing the delivery fee and reducing the opening product bid price by the same amount, generated higher end prices in the auctions. However, once prices became inclusive of delivery fees, only revenue for silver coins increases

using that strategy. Changing the delivery cost also has a negligible impact on the number of bidders.

Huck and Wallace (2015) conduct a laboratory experiment that focuses explicitly on the impact of framing effects on consumer choice. In contrast with Morgan et al. (2006), Huck and Wallace (2015) focus on the responses of consumers to different framing strategies, rather than automating the demand side of the market. Instead, the supply side of their market is computer automated. They compare the optimality of participants' purchasing and search behaviour under a simple linear per unit tariff with behaviour under five common retail framing effects: Reference pricing, drip pricing, time-limited offers, complex/sale pricing and baiting. Search behaviour refers to the order and amount of search that participants undertake. Purchasing decisions refer to how much of the product the subjects purchase in each period.

At the beginning of each period, participants receive information regarding the four products that they may be asked to buy, including the range of prices for each product, from which a price will be drawn randomly for each of the two fictional sellers. Each participant must then choose a store to visit, which in several treatments is a random decision. Visiting a store incurs a cost to the participant, which reduces their score. Upon visiting the store, the buyer learns the price of the product at that store and decides how much to buy. Each buyer can purchase a maximum of four units of the product in each period. If four units are not purchased from the first seller, the buyer can return to the home screen.

Each buyer can subsequently visit the competitor to learn their price or return to the first searched store, which incurs a second unit of the search cost. Again, the player chooses how much to buy before returning to the home screen. Any additional search beyond two firms incurs additional units of the search cost. Participants act independently and products cannot be returned. The participant's objective is to maximise their score by purchasing four units at the lowest price.

The benchmark is a linear tariff that contains a fixed per unit price of the product. Under reference pricing, a 'was' price is displayed above the current price, in addition to the associated discount that the current price constitutes. Participants know that the *sellers* have not operated before. Drip pricing partitions the price into three components. The consumer first learns a base price, to which a further charge of 5%-15% of the base price is added if the consumer attempts to buy. If the consumer continues, a second fee is levied, which is 10%-20% of the base price. The consumer can decide not to buy when new fees are added.

Time-limited offers are emulated by imposing a 1 period time constraint on the price of the first searched firm. If the consumer returns to this firm at a later time in the period, the price will have changed. Notably, the price is a new random draw and may increase or decrease. A treatment for complex pricing is created using a non-linear pricing structure with a simple '3 for

2' promotion. Finally, baiting involves advertising prices on the homepage before the consumer decides which firm to search. However, the advertised prices are subject to stock limitations and may not be available when the consumer visits the store.

In total, 166 students play the role of buyers. Participants are given the base treatment without any framing and two other randomly selected treatment effects. The game is replayed for 30 periods, distributed evenly between the three treatment effects, which are presented randomly. This produces nearly 5000 observations of search and buying behaviour. Four products are also used but only one of the products is available in each period. The products vary only in their price ranges and their search cost.

In all cases, the pricing strategy followed by the two sellers is symmetric: If one firm offers a multi-buy deal, the competitor follows the same framing effect. It would, therefore, be interesting to analyse how the outcomes for firms and consumers change when firms follow different promotion strategies. For example, how do consumers behave if one firm places a time constraint whilst the rival uses drip pricing? This is one specific opportunity for future research that would draw the experimental design closer to real world retail markets.

The participants' actual search and purchasing strategies in the experiment are compared with the theoretically optimal strategy for that period. The change in optimality due to framing, however, is measured with respect to their performance in the benchmark treatment. A probit regression is conducted, where the amount of errors are regressed on: The prices of the products, the value of the search cost, a time-trend variable to account for learning and indicator variables to account for which frame was applied.

Their main result is that individuals are highly sensitive to framing. Three of the five treatments generate a statistically significant loss for consumers, relative to the baseline treatment. Drip pricing induces the most pronounced loss, followed by time-limited offers and baiting. Using a '3 for 2' offer and a fictional previous listing price did not yield statistically meaningful changes in consumer surplus. The error rate is higher under all framing treatments than the linear pricing schedule. However, decomposing the errors into search and purchasing mistakes, Huck & Wallace (2015) show that search patterns are more significantly affected by framing than the decision of how much to buy. Time-limited pricing, drip pricing, baiting and reference pricing all lead to a significant increase in the amount of search errors but only drip pricing and baiting stimulate the purchasing error rate.

Surprisingly, introducing time-sensitive prices yields the highest increase in the error rate, with search errors increasing by 20%. Drip pricing, baiting and reference pricing yield 11%, 9% and 5% increases in search errors, respectively. Buyers also undertake excessive search in the baseline treatment but framing induces consumers to buy sooner by reducing the amount of search. For reference pricing, the framing-induced reduction in search is approximately equal to the initial

over-search but for all other treatments, the reduction in search due by framing exceeds the initial tendency to over-search.

The impact of search costs on participant's behaviour, controlled by varying the reduction in points that participants incur for visiting a firm, is multi-directional. Increasing the cost of search does not influence the purchasing error rate but does reduce the rate of search errors in three out of the five treatments. The largest reduction in search errors occurs under time-limited offers (-2.52), followed by baiting (-2.25) and drip pricing (-1.74).

Considering the artificial supply side of this market under the baseline treatment, the first searched shop sells 12% less of its product than would occur if consumers followed the optimal strategy but the second searched shop outsells its expected volume by 7%. When framing is introduced, consumers search less. As a consequence, the first searched firm sells significantly more under drip pricing, baiting and complex pricing and exceeds the volume that would be expected under optimal consumer behaviour. The second searched firm now sells less than anticipated.

Time-limited offers yield an anomaly because the firm that is searched second by participants sells 68.5% more than predicted. The intuition follows that consumers who leave the first store never consider returning because they anticipate a higher price. Whilst this is not necessarily the case, the fact that they do not return implies that they will never learn this. Participants also become better decision-makers over time and individuals that score higher on an IQ test tend to follow the optimal choice procedure more closely. However, interacting IQ with time-trend effects indicates that learning is particularly slow for time-limited offers because consumers never return to the first searched store.

2.3.2.1 Differentiated Products

In many markets, products are genuinely differentiated. Perhaps, most obviously, in terms of quality. Therefore, when consumers consider different offers, they must identify the lowest quality-adjusted price. This introduces a complication for consumers that is not present in Hossain & Morgan (2006), which is the starting point for Kalayci and Potters (2011).

They study an experiment where products are vertically differentiated and both buyers and sellers are played by human participants. They analyse how the quality of a seller's product influences pricing and obfuscation decisions. In the first stage of each period, two sellers choose the complexity of their pricing structure, conditional on the quality of their product relative to their competitor. Complexity is interpreted as the number of attributes over which the quality of a product is distributed. If the seller chooses one attribute, all of the quality value is allocated to dimension 1. If the seller chooses two dimensions, the total quality value is distributed randomly across dimensions 1 and 2. The same pattern continues for up to 5 available attributes.

In the second stage of each period, obfuscation decisions become common knowledge and sellers choose their price. In stage three, consumers have 15 seconds to decide which product, if any, to purchase. Buyers maximise their payoff:

$$\text{Payoff} = 5 \cdot q_5 + 4 \cdot q_4 + 3 \cdot q_3 + 2 \cdot q_2 + 1 \cdot q_1 - p \quad (2.19)$$

q_i denotes the quality of the i 'th attribute of each product and p denotes the price. Therefore, to maximise their earnings, buyers must calculate the weight-adjusted valuation of each seller's product minus their price. Prices and quality are presented to buyers in the following format:

Product/Weight	5	4	3	2	1	Price
Good A	0	4	4	11	36	50
Good B	0	0	0	0	73	45

Figure 2.4: Kalayci and Potters (2011).

In this example, seller A has chosen 4 attributes, whilst seller B has chosen only 1. Consumers are expected to make more mistakes than if seller A had used one dimension.

At the end of each period, buyers and sellers review their performance. For sellers this includes the price, quality, number of dimensions, sales and profit for both firms (Kalayci and Potters, 2011). For buyers this includes the prices of the products in the market, their purchases and their payoff. Buyers are not informed of the underlying incentive structure for the sellers but they do know that sellers will choose the number of components. In total, 48 students participate and each student is randomly assigned to the position of buyer or seller.

In theory, there exists two potential strategies for firms, depending on the magnitude of the quality difference between firms, relative to the rate at which using more dimensions confuses buyers. Either, both the low and high quality firm may find it profitable to obfuscate as much as possible, or only the low quality seller may choose to obfuscate. The low quality seller always obfuscates because informed consumers will never buy from him. The strategy for the higher quality seller is less clear cut. If the noise that the rival creates is large, relative to the quality difference, firm 1 will find it profitable to obfuscate. However, if the noise created by the rival is relatively small, compared to the differences in quality, firm 1 may not obfuscate.

In their experiment, sellers use 2.6 attributes on average and the number of attributes declines marginally throughout the experiment. When at least one seller opts for multiple attributes, consumers make 35% more mistakes.¹⁵ However, there is no marginal effect on the rate of mistakes for a firm adopting a three, four or five part tariff, instead of a two-part tariff. The probability of

¹⁵A mistake is defined as the buyer choosing the product with the lower payoff or not purchasing at all.

a mistake also increases when the firm offering the highest consumer payoff uses more attributes than the rival, which lends support to buyers preferring simple tariff structures. This is discussed further in Section 2.4. Prices are also higher when the number of attributes increases.

Returning to the impact of vertical differentiation, Kalayci and Potters (2011) find that low quality firms obfuscate more. On average, a multi-part tariff is applied with 65% probability and 5% less for higher quality sellers. This difference is relatively small, which implies that the effect of obfuscation on consumer choice is large, relative to the differences in quality. Kalayci and Potters (2011) repeat their study with computerised buyers that always purchase from the firm offering the highest surplus. They find that prices are now lower, which confirms that sellers combine higher prices with more attributes because they expect confused consumers to make more mistakes.

In a related field experiment, Bertrand et al. (2010) study the effects of pricing and advertising strategies on the demand for loans sold by a lender in South Africa. The lender distributed promotional leaflets to 53,194 individuals who had borrowed from the firm in the preceding 6-24 months. The interest rates, information content and promotional materials contained on each advert were varied according to psychological principles and commercial interests to detect the effects of particular advertising and pricing methods. The clients are typically the working poor who borrow small amounts for a short period of time, with an Annual Percentage Rate in the range of 200%.

Adverts were tailored to the risk profile of the individual, which enabled the lender to vary the rate of interest across the adverts and estimate price elasticities, whilst incorporating the probability of default. The adverts were also pre-approved and a customer needed only to telephone or visit the local branch to accept the loan. Consumers are also accustomed to receiving monthly product information but this firm had never previously offered any promotions.

A wide range of treatments were applied across the adverts. Firstly, a photograph of an attractive individual was placed on 80% of the adverts, with equal weight placed on males and females and a bias in favour of the clients race. Second, each advert contained either one or four numerical examples of the cost of a fictional loan, which were relevant for the individual based on their previous borrowing. All examples included the total loan amount and the maturity, but only some adverts displayed the interest rate explicitly, which acts as a third treatment.

The fourth treatment stated possible uses for the cash, which varied from suggesting a particular use to stating 'anything you want' (Bertrand et al., 2010). Some adverts also highlighted the price of the same loan offered by a competitor in the market, which is the fifth treatment. The sixth treatment is a promotional raffle where customers who use the loan are entered into a draw to win a mobile phone. The seventh treatment emphasised the lender's ability to speak the same language as the customer. The eighth treatment highlighted on some of the adverts that the

discounted rate was temporary. The final class of treatments varied the length of the promotion between 2 - 6 weeks, which allows an examination of time-sensitive offers.

Their main results are as follows. The probability that a consumer purchases the loan increases by 4% for every 13% reduction in the interest rate, generating a price elasticity of approximately -0.28. Increasing the rate of interest also reduces the amount that individuals elect to borrow, with an elasticity of -0.34. They show that the promotion generated additional demand for loans, rather than causing borrowers to substitute from other providers.

Only three of the eight treatments stimulate demand. These are: Providing no examples of the potential use of the loan, providing one example loan calculation for the cost of a loan and including a female picture on the posted advert. Following one of these treatments can increase demand by the same amount as reducing interest rates by 25%. However, Bertrand et al (2010) uncover a significant difference in the effects of treatments across genders. The three significant treatment effects increased demand from male clients but had no bearing on the loan decisions of females. Notably, comparing the rate of interest to their competitors, presenting additional numerical examples and emphasising the special price did not affect demand (Bertrand et al., 2010). Relaxing the time constraint of the promotion also stimulates demand: A two week extension increases total demand by 3%. This figure is particularly large given that the average loan uptake is 8.5%.

In a follow-up article, Kalayci (2015) reports several laboratory experiments that investigate the relationship between prices and price complexity in a duopoly market with homogeneous products. Much of the groundwork for this experiment draws from Kalayci & Potters (2011), with two main differences. Firstly, quality is now symmetric across sellers and exists only for buyers to value the product. Secondly, the role of buyers and sellers is now played by both human participants and computerised buyers, depending on the treatment. The two sellers compete in prices and price complexity. This allows us to zoom in on the interaction between the price that a seller charges and the obfuscation they use.

The theoretical framework draws on Carlin (2009) in which sellers simultaneously select their price and its complexity.¹⁶ Kalayci (2015) follows their earlier study (Kalayci and Potters, 2011) by interpreting complexity as the number of partitions that the total price is distributed across. Using more partitions increases the difficulty of comparisons for buyers. Sellers choose the total price and the number of components but not how the price is distributed across the different dimensions, which is an interesting opportunity for future work.

Two main hypotheses of Carlin (2009) are tested. Firstly, firms that charge higher prices should choose higher complexity. Secondly, complexity serves as neither a strategic complement or substitute for sellers, which implies that obfuscation decisions are independent. Three experiments

¹⁶See Section 2.2.1.1.

are reported, with different treatments for each. In all cases, sellers choose the price and the number of partitions for the price in the first stage and up to three partitions are allowed. The buyer must subsequently choose, within a specified time frame, which of the two sellers to purchase from (if any). Their objective is to maximise a weighted sum payoff function:

$$\text{Payoff} = \text{Quality} - (\text{Fee } 1 + 2 \cdot \text{Fee } 2 + 3 \cdot \text{Fee } 3) \quad (2.20)$$

Similar to the 2011 study, both players receive a feedback report at the end of each period and each participant plays for 30 periods.¹⁷

The main results are delivered in two of the experiments, which feature two and three treatments, respectively. In the first study, human sellers compete for human buyers. The ability of buyers is controlled by adjusting the time available to make a decision. In treatment one, buyers have 10 seconds to choose. In treatment two, buyers have 45 seconds to choose. In the second study, the 10 second treatment is repeated because the location of the experiment is changed. This is referred to as the *human* treatment. In treatments two and three, human buyers are replaced by computers. In treatment two, buyers are programmed as fully rational agents who buy from the cheapest seller. In treatment three, the buyers are programmed to behave as boundedly rational consumers. This involves 5% of the buyers making mistakes when firms choose single unit pricing and this fraction increases by 5% for every additional partition used by any firm, to a maximum of 35%. In total, this produces 5 treatments: *10 seconds*, *45 seconds*, *human*, *rational* and *boundedly rational*.

The first experiment indicates that a seller with a higher price is more likely to use three partitions (32%) than a firm with a lower price (12%) in the *10 second* treatment but this correlation disappears when buyers have more time in the *45 second* treatment. However, when the same 10 second experiment (*Human*) is re-run in the second study, no correlation is present. The positive association is restored in the market with *boundedly rational* robot buyers, with more expensive sellers using three dimensions in 49% of periods and less expensive sellers using three dimensions only half as frequently. Surprisingly, average prices are not statistically different between the 10 and 45 second treatment. Therefore, providing consumers with more time does not compete away firms' profits.

In the second experiment, average prices are lowest in the robotic *rational* buyer treatment and highest when robots act as *boundedly rational* buyers. Average prices also decline throughout all of the experiments, except for the *boundedly rational* robot treatment where prices begin to decline before quickly climbing above their initial level.

¹⁷At the feedback stage, buyers learn the cost of their purchases and their payoff but they do not learn the counterfactual payoff if they had purchased from the other supplier. Sellers learn an array of market information including the price, quality, price complexity, sales and profit of both sellers, including information from previous periods.

Kalayci (2015) offers two possible explanations for the conflicting results regarding the correlation between price and complexity. Firstly, the assumption of complexity not being a substitute amongst sellers may be misplaced. This implies that the marginal effect of complexity for a firm will be reduced when another firm also selects high complexity. Secondly, the amount of mistakes is negatively related to the price difference in the human treatment but independent in the boundedly rational robot treatment, which suggests that price dispersion plays a role.

To address this, a third experiment is executed in which human participants play the role of buyers over an extended 60 period game with 10 seconds of decision-making time per period but the role of sellers is automated by a computer. This allows price dispersion to be controlled directly. The computer is programmed for three treatments: Both sellers use three price components, both sellers use one price component and firm 1 uses one component but firm 2 uses three components.

For 50% of the third experiment, quality and price are drawn randomly from [60, 100] and [20, 60], respectively. For the remaining 50%, quality and price are scaled to the range [600, 1000] and [200, 600], respectively. Prices are also independent in half of the treatments and constrained to a 5 unit dispersion (50 unit dispersion for the scaled case) in the other half. The payoff structure for buyers also changes. Specifically, a subset of participants play for a fixed reward for every period in which they buy optimally and nothing otherwise.

Kalayci (2015) shows that the absolute differences in prices influences the amount of mistakes buyers make. However, buyers make fewer errors when prices are low and this price level effect is amplified when participants are paid based on the variable payoff, rather than the fixed binary fee. The absolute quantity of mistakes, however, is invariant to the specification of the buyer's payoff. The amount of mistakes is also minimised when both sellers use a single price but there is no discernible difference in the amount of mistakes between the treatments where one or both sellers use the three part tariff.

These results lend partial support to the idea that more expensive firms will obfuscate more. In particular, when the experiment is calibrated with computerised buyers that adhere to Carlin's (2009) assumptions, his equilibrium predictions are correct. However, in practice, price is not the sole determinant of obfuscation levels and the observation that many sellers use two dimensions, instead of alternating between one and three, requires further investigation.

Gu and Wenzel (2015) consider a different asymmetry across firms to explore how more prominent sellers obfuscate and price their product, compared to less prominent firms. In their theoretical model (Gu and Wenzel, 2014), they consider a duopoly of firms that sell to a population of Varian (1980) style consumers: Some consumers are informed and buy at the lowest price but others are uninformed and cannot compare prices. However, instead of uninformed consumers buying randomly, they suppose that an exogenously determined larger fixed fraction will buy from one

of the sellers, to capture market leadership effects. The division of uninformed consumers is also known to sellers.

At stage 1, firms select a scalar obfuscation level, which aggregates to market complexity and determines the proportion of informed consumers, in the style of Carlin (2009). When aggregate obfuscation increases, fewer consumers are informed. At stage 2, obfuscation decisions become common knowledge and firms set their price. They show that the firm that is guaranteed a larger share of uninformed consumers selects maximum obfuscation and charges higher prices than the rival, who may select below maximum obfuscation depending on the proportion of consumers that are confused by the obfuscation of the prominent seller. That is, there exists a specific fraction of informed consumers at which the less prominent firm maximises profit. Therefore, policies that reduce the effectiveness of obfuscation induce a higher level of obfuscation by the rival, to compensate for the increase in the fraction of informed buyers.

In their experiment, the timing is identical but human participants play the role of sellers and the supply side of the market is automated by a computer. At the beginning of each period, each seller is informed whether they are the 'prominent' or 'non-prominent' firm in a randomly matched duopoly market. In stage 1, sellers make a binary decision between obfuscating and not obfuscating, where obfuscation reduces the proportion of informed consumers by 20%. In stage 2, stage 1 decisions become common knowledge and sellers select prices from the interval $[0, 100]$ (Gu and Wenzel, 2015). At the end of each period, sellers learn their competitor's price and their profits. The game is repeated for 25 periods and in each period the seller faces a different competitor.

They show that the firm that is guaranteed more of the confused population, obfuscates more, which is consistent with their theoretical model. However, when the fraction of consumers that become confused for each firm that obfuscates decreases, the non-prominent seller increases his obfuscation by a smaller quantity than predicted. In addition, when the 'prominence' of the market leader is less pronounced, as captured by a reduction in the sharing rule of uninformed consumers from 90% to 60% for the 'prominent' seller, firms obfuscate more and charge higher prices. Therefore, asymmetries between firms can influence both obfuscation decisions and the efficacy of policies intended to limit the effects of obfuscation.

2.3.2.2 Natural Increases in Choice

In addition to obfuscation motives for firms to increase the choice available to consumers, the range of options available in a market can also increase for competitive or policy reasons. Hastings and Tejada-Ashton (2008) study the effects of liberalising the Mexican retirement investment industry in 1997, which was motivated by a desire to improve market efficiency and introduce greater choice for savers. In Mexico there exists no default or state promoted investment schemes.

All workers must select their own investment plan. Switching is relatively straight-forward with the average person switching once every three years (Hastings and Tejada-Ashton, 2008).

Investment products are also tightly regulated and financial institutions must offer a minimum of two government specified funds that differ in their risk profile. Consumers are highly engaged and aware of the switching opportunities available but many choose inefficient investment plans. The average saver invests in the 12th most efficient fund, out of the 17 available, and overpays by twice the price of the cheapest provider.

Hastings & Tejada-Ashton (2013) use survey and experimental data to understand how the framing of fees can influence consumers and analyse the impact of socio-economic factors on decision-making, which is an under-researched topic. In total, 763 individuals participate in the four stage survey. Firstly, participants provide detailed personal information on their background. Secondly, individuals report their current retirement investments and their knowledge of the system. Thirdly, a series of finance questions gauge financial literacy. They find that individuals with higher income and educational attainment, perform better. Finally, participants select three investment funds that they would recommend to a family member, based on a fictional financial profile.

The sample are representative of the Mexican population, except low income consumers are under-represented. Participants are also reasonably aware of their investment plans, with 81% able to state their current retirement fund. The experiment tests several framing effects, including different fee structures and different periods of time for reporting the incurred fees (i.e. annually and every 10 years).

They find that participants who perform highest in the financial literacy assessment are more sensitive to the fees charged on an account when recommending an investment scheme. However, when the structure of the fees is marginally simplified, by denominating the fees in the domestic currency rather than a percentage charge of the fund size, less literate investors became significantly more price sensitive. More specifically, restating prices in Pesos, instead of percentages, yields a 25% - 55% increase in price sensitivity across financially illiterate consumers, with no impact on financially literate investors.

By comparison, policy initiatives that educate consumers, generate an estimated 75% - 134% increase in the price sensitivity of investors. Therefore, reasonably simple and inexpensive changes to information formats should be mandated before expensive education initiatives are introduced (Hastings & Tejada-Ashton, 2008). Their data also indicates that females and individuals with lower salaries rely heavily on guidance from their employer when selecting a pension plan. In contrast, males, who represent a larger share of the working population in Mexico, seek more guidance from the regulator. Overall, lower income and less educated consumers pay less attention to account fees and focus on non-price factors, including access to in-store support. These

results are salient because, in contrast with the prevailing focus on aggregate consumer welfare, Hastings & Tejeda-Ashton (2008) demonstrate that policies that simplify the framing of product information will generate a larger benefit for less educated and lower income individuals.

2.3.2.3 Saliency

In a forthcoming publication, Dertwinkel-Kalt et al. (2016) provide the first experimental examination of saliency theory (Bordalo et al., 2013). In contrast with much of the experimental literature, the 169 participants in this study were not asked to complete any tasks. Instead, participants were asked to spend 45 minutes in front of a computer, for which they would receive a fixed amount of money. Before the 45 minutes begins, individuals are offered the opportunity to buy either high or low speed internet to use during the period. Their computer usage is unmonitored and no other technology or communication devices are permitted. Therefore, their choice of internet speed determines how fast they can browse online. The price of internet is deducted from their participation fee. They study how the proportion of individuals who buy faster internet changes as prices change and individual's expectations of the prices change.

Saliency theory predicts that a uniform increase in the prices of two vertically differentiated products should shift demand towards the higher quality item. When prices are marked up, quality differences become salient and price differences become less salient. In Section 2.2.3 we discussed how demand for wine shifts towards higher quality bottles when prices are marked up in a restaurant, compared to a supermarket (Bordalo et al., 2013). Saliency theory also predicts that individuals are more likely to purchase the high quality product when prices are expected to be high than if prices are expected to be low. If the price of wine is expected to be high in the restaurant, individuals are more likely to buy it (Bordalo et al., 2013).

Participants are randomly split between three treatments. The first two treatment groups, *Low Price* (LP) and *High Price* (HP), are given the internet prices a few days prior to the experiment. Prices in LP are £0.50 and £1.50 for low and high speed internet, respectively, and the participation fee is £12. Prices in HP are £3 higher at £3.50 and £4.50, for low and high speed internet, respectively. The participation fee is also increased in HP to £15 to negate any income effects. Comparing the demand allocations between the high and low price treatments isolates the impact of a uniform increase in prices.

To test the role of expectations, a third treatment: *Unexpected High Price* (UHP), is constructed. In this session, participants receive the price information a few days in advance but they are informed that the prices may follow either the LP or HP treatment prices. On the day, the prices and show-up fee follow the HP treatment and comparing HP and UHP allows us to examine the impact of expectations on individual choice.

Their results indicate strong support for salience theory, which cannot be systematically explained by any other theory of individual choice.¹⁸ When prices are low (LP), only 28.1% of the group choose high speed internet. When prices are uniformly higher (HP), 45.8% choose high speed internet. Therefore, when prices are marked up by a fixed amount, individuals inflate the relative attention given to quality and reduce the relative attention given to prices. When prices are uncertain for consumers (UHP), only 26.4% of individuals pay for high speed internet. This is significantly lower than the 45.8% that upgrade when prices are known with certainty. Therefore, prices become more salient to consumers when they are unexpectedly high, than when they are expectedly high (Dertwinkel-Kalt et al., 2016).

¹⁸See Dertwinkel-Kalt et al. (2016) for a review of related alternate theories of choice, including Prospect Theory (Kahneman and Tversky, 1979), Personal Equilibrium (Kőszegi and Rabin, 2006), Focusing Theory (Kőszegi and Szeidl, 2013), Relative Thinking, and Comparisons and Choice.

2.4 Simplicity Biases: Theory & Evidence

One of the most frequent assumptions within the pricing and obfuscation literature is that consumers who cannot understand prices buy from a randomly chosen seller (possibly with an exogenous bias towards a market leader) (Carlin, 2009; Chioveanu and Zhou, 2013; Eliaz and Spiegler, 2011a; Gu and Wenzel, 2014; Narasimhan, 1988; Piccione and Spiegler, 2012; Varian, 1980). However, experimental and empirical evidence consistently challenges this assumption because consumers are complexity averse. Therefore, when confronted with a range of simple and complex offers, consumers are biased towards simple options (Sonsino et al., 2002; Homburg et al., 2014). This section reviews the evidence of complexity aversion and motivates its inclusion in models of competitive obfuscation, which is a principal topic of this thesis.

Some of the earliest reports of complexity aversion were detected in experimental appraisals of expected utility theory. Notably, the experiments of Huck and Weizsäcker (1999), Mador et al. (2000) and Sonsino et al. (2002) gave traction to the idea that individuals will discriminate against complex alternatives. These studies analyse how individuals choose between different lotteries, which in its simplest form is an uncertain future income stream that assigns probabilities to possible future payoffs or alternatives. By varying the quantity of alternatives or the intrinsic complexity of the alternatives, we can isolate how individuals respond to complexity.¹⁹

Huck and Weizsäcker (1999) ask 71 students to choose between two single-period lotteries that vary in: The number of possible outcomes, the number of digits used to express the probabilities and outcomes, the magnitude of the payoffs and the standard deviation of the payoff alternatives. There is no time constraint and students receive their lottery payouts at the end of the experiment. They find that individuals systematically depart from expected value maximisation in a way that cannot be explained by risk aversion. When students face two lotteries with an equal quantity of alternatives, increasing the amount of alternatives causes more deviations away from expected value theory. If agents were fully rational and informed, participants' choices would be independent of the number of alternatives (Huck and Weizsäcker, 1999). Moreover, when one lottery contains fewer alternatives than the other, pupils are consistently biased towards the less complex lottery.

Cluster analysis also reveals that the strongest simplicity bias is detected amongst participants whose decisions are not distinguishable from pure guesses. These pupils achieve the lowest grades in academic mathematics courses, which implies that the preference for simplicity is most pronounced for individuals who cannot evaluate the available options. Therefore, obfuscation has distributional implications across consumers. This reinforces the concerns of Hastings and Tejeda-Ashton (2008).

¹⁹Bruce and Johnson (1996), amongst others, focus on how increasing the complexity of the choice problem itself can affect individual behaviour, rather than the complexity of the alternatives.

In a closely related study, Mador et al. (2000) analyse the role of complexity in multi-period lotteries that assign payoff streams to future payoffs. Therefore, the value of a multi-period lottery is the discount-adjusted sum of the probability-weighted payoffs. They test whether individuals choose the utility maximising alternative, conditional on future payouts being valued less than payouts made today, which is a core prediction of discounted expected utility theory.

Defining complexity in multi-period lotteries is more involved. Firstly, lotteries with more alternatives are more complex. Secondly, lotteries with more dispersed numerical values are more complex. Thirdly, dispersion in the probabilities assigned to the possible outcomes also matter. Mador et al. (2000) construct a complexity evaluation structure that defines the most complex lottery in the following way. Let the two lotteries be denoted L_1 and L_2 , with $L_i = (p_i, X_i)$. p_i denotes the vector of probabilities assigned to each possible alternative of lottery i and X_i denotes the payoff matrix with dimensions $m \cdot T$, such that with probability p_j , the T period payoff stream will be the stream given by the j^{th} row of matrix X (Mador et al., 2000). Lottery L_1 is more complex than lottery L_2 if: (i) X_1 contains more or equal rows or columns than X_2 , (ii) X_1 contains more unique elements than X_2 , or (iii) p_1 contains more or the same number of different components as p_2 .

73 undergraduate students with basic training in economics were asked to bid in prices for nine lotteries, rather than choosing between different lotteries. The beauty of asking subjects to price the lotteries is that we can observe their relative preferences for the different alternatives. If a participant bids a lower price, they must value the lottery less. At the beginning of the experiment, students receive a cash endowment with which to buy lotteries and students are remunerated with a minimum show-up fee plus the payoff from any lotteries for which they win, by bidding the highest. The lottery payoff is calculated as the lottery outcome minus the price bid.

In the first four lotteries, presented randomly to subjects, the lotteries are increasing in complexity. The discounted expected utility for lottery one is lower than any of the other lotteries and should receive the lowest bid price from students. However, 28%, 42% and 62% of students bid a lower price for lotteries two, three and four than lottery one, respectively. Average bid prices also decline as the complexity of the lottery increases and lottery four stochastically dominates lottery one, but the average bid for lottery 1 was significantly higher. Therefore, individuals exhibit a lower willingness to pay for more complicated alternatives.

Lotteries five and six are personalised to the individual's offer for lottery four. Specifically, in lottery five participants choose between a fixed sum today or the price that they offered for lottery four, paid in three months. Lottery six involves choosing between a slightly larger fixed sum today or the outcome of the fourth lottery in three months. If the individual did not offer to pay more for lottery four than the certainty equivalent of these lotteries, theory predicts that

subjects should offer to pay more for lottery six but 45% do not. Their explanation is that lottery six is more complicated than lottery five, which causes individuals to reduce their bids.

Lotteries seven and eight are relatively simple. Lottery seven pays today whilst lottery eight pays in three months. Lottery nine is the bundle of lotteries seven and eight together. Taking account of risk aversion, individuals should bid more for lottery nine than the sum of their bids for lotteries seven and eight, but only 30% do. The intuition is that the combination lottery is more complex for individuals to evaluate, which leads to lower bid prices.

Sonsino et al. (2002) present an experiment with lotteries that span multiple periods but participants choose between lotteries, rather than pricing them. They conduct in-class experiments with 97 undergraduate and graduate students that have completed basic courses in probability and economics. Students complete 14 tasks over three sessions, with 1-2 week intervals between meetings. Subjects are incentivised to select their preferred lottery using deferred cheques, which pay out at six month intervals for each deferred payoff period in the experiment. The actual payout that students receive is also proportional to the payoff earned during the tests.²⁰ Similar to Mador et al. (2000), their finite multi-period lottery is defined as (p, X) . Complexity is controlled by adjusting the dimensions of X , such that increasing the number of periods or the number of alternatives associated with the lottery, makes the lottery more complex. Mathematically, the complexity level is rated as the product of the two dimensions.²¹

Three axioms are tested. Firstly, convexity, which states that if two lotteries are preferred to a third, the decision maker should prefer any convex combination of the two preferred lotteries to the third. Secondly, stochastic dominance, which states that a lottery that pays out more in every period should be preferred. Thirdly, independence of irrelevant periods, which states that if the distributions of the payoffs in two lotteries are identical except in the payoff for one period, the lottery with the highest payout on that asymmetric period should be preferred.

In each task, subjects select one of two lotteries. Up to six variants are created for each task, which randomise the presentation of the information and the alternatives. In the first four tasks, alternative A is the same throughout but alternative B varies. Alternative B in task three is a convex combination of task 1. Therefore, individuals who select alternative B in task 1 should stay with alternative B in task 3. Alternative B in task 4 is a marginal improvement on alternative B in task 3. Therefore, participants who select B in task 3, should continue with B in task 4. However, the results do not adhere with this theory. Instead, the proportion of individuals that select alternative

²⁰The tasks of the third session were also given to a further group of 23 students, to ensure the absence of wealth effects throughout the experiment. This is confirmed.

²¹For example, a lottery that is completed in one period with two possible alternatives has a complexity rating of 2. A lottery that develops over three periods with four possible alternatives in each period has a complexity rating of 12. The maximum number of periods in the experiment is three.

B decreases as the four tasks progress. In particular, 17 individuals choose alternative B in task 3 but switch to alternative A in task 4 and only 7 switch in the other direction.

Sonsino et al. (2002) argue that complexity aversion can reconcile this behaviour. Alternative B in the fourth lottery is more complicated than in lottery 3 because more potential outcomes are included. B is also less complicated in tasks 1 and 2, which leads more individuals to select B than in the subsequent tasks.

In the second session, participants complete tasks 5-7, which involve lotteries with two periods of payoffs. Alternative A is fixed throughout. The payoffs of B in task 6 stochastically dominate the payoffs of B in task 5. Therefore, students that select B in lottery 5 should not switch to A in lottery 6. Alternative B in lottery 7 stochastically dominates alternative B in lottery 6, such that individuals should not choose option B in lottery 6, without choosing alternative B in lottery 7. Their results partially reaffirm their previous findings as subjects switch from B in lottery 5 to A in lottery 6. However, participants do not switch away from B in lottery 7. They show in subsequent econometric analysis that the increase in complexity creates additional variance in individuals' choice procedures, which counteracts the complexity effect.

In tasks 8-10, both alternative A and B vary to test the independence of irrelevant periods. They find strong evidence against both independence of irrelevant periods and discounted expected utility theory. This strengthens the idea that complexity aversion interacts with other cognitive constraints (Sonsino et al., 2002).

Task 11 is a choice between a lottery (A) and the certain payoff of the expected value of the lottery (B). If individuals are risk averse they should prefer the expected value but 30% choose the lottery. In task 12, A is the same as in lottery 11 but B differs. Instead of offering the expected value of A with certainty, B is a return ($R(X)$) from a financial security (X). Specifically, B comprises two equally weighted and independent copies of the security, with a return: $\frac{1}{2}R(X_1) + \frac{1}{2}R(X_2)$ where X_i denotes the i^{th} independent copy of X (Sonsino et al., 2002). If participants follow expected utility theory and are risk averse, students who select B in task 11 should also select B in task 12. However, only 53% choose alternative B, which is a statistically significant reduction from 69% in task 11.

One explanation is that risk averse individuals are perturbed by the complexity of alternative B in task 12. However, conditional on selecting alternative A in task 11, 52% of subjects choose alternative B in task 12, which cannot be explained by complexity aversion alone. To reconcile this, Sonsino et al (2002) argue that increasing complexity has two effects. Firstly, as alternatives become more complex, decision making becomes less accurate. This causes some participants to switch from A to B or B to A between tasks 11 and 12. Secondly, when an alternative becomes more complex, individuals discriminate against it, which can explain why individuals switch from

B to A between tasks 11 and 12. Therefore, these two effects will interactively determine subjects' decisions.

Task 13 analyses subjects' tastes for diversification. B is a diversified equivalent to alternative A. 69% of individuals displayed risk aversion in task 11. Therefore, a larger fraction should prefer B in task 13 but only 40% do. In task 14, the alternatives are identical except the framing of the lottery in alternative B changes from two independent lotteries to one two-period lottery (Sonsino et al., 2002). In addition, payoffs for alternative B are marginally increased to obtain stochastic dominance. Discounted expected utility theory suggests that subjects should prefer the diversified option but only two thirds do. They argue that alternative A appears less complex, which drives this result.

Sonsino et al. (2002) take an additional step by applying formal econometric apparatus to verify the statistical significance of their complexity aversion effect. This also enables complexity aversion to be disentangled from the additional noise-creating effect that complexity has for decision-makers. Two modelling approaches are outlined: Differential and proportional. In the differential approach, the *value adjusted to complexity* (VAC) of each alternative (A_j) for individual i is written:

$$VAC_i(A_j) = \beta_i^0 + \beta_i^V V(A_j) + \beta_i^C COMP(A_j) + \eta_i(A_j) \quad (2.21)$$

$V(A_j)$ is the value of the lottery, independent of its complexity level. $COMP(A_j)$ is the complexity rating of alternative j . $\eta_i(A_j)$ is the standard error. The model implies that alternative A should be selected if and only if: $VAC_i(A_j) > VAC_i(B_j)$. In the proportional approach:

$$VAC_i(A_j) = \beta_i^0 V(A_j)^{\beta_i^V} COMP(A_j)^{\beta_i^C} e^{\eta_i(A_j)} \quad (2.22)$$

Alternative A_j is preferred if and only if: $\frac{VAC_i(A_j)}{VAC_i(B_j)} > 1$.²²

Fixed effects and random effects models are estimated, to account for heterogeneity across the participants.²³ They find strong evidence that increasing the complexity of an alternative causes individuals to encounter greater noise in their decision-making as the variance of the error term increases. Increasing the expected value of a lottery (independent of complexity) also increases the probability that this alternative will be selected. However, more saliently, increasing the coefficient for the complexity of an alternative causes a statistically significant reduction in the probability that the alternative will be chosen.

In a subsequent experiment, Sonsino et al. (2002) also show that the bias towards simple alternatives survives when the expected value of a lottery is displayed alongside the alternatives that participants face. As a result, they propose that a complexity adjusted model of discounted expected utility may have greater explanatory power in modelling individual choice behaviour.

²²The interested reader is referred to the article for further derivations and mathematical intricacies of the different modelling approaches, which for brevity are not repeated in this review.

²³Fixed effects imply: $\beta_i^V = \beta^V$ and $\beta_i^C = \beta^C$ for every individual i (Sonsino et al., 2002).

One final contribution of Sonsino et al. (2002) is that combining their results with the previous literature on choice under uncertainty, particularly the work of Prelec and Loewenstein (1991), suggests that the marginal effect of complexity may be decreasing. Therefore, initial increases in the complexity of alternatives should have a larger impact on consumer choice than subsequent increases in complexity. This complements the results in Section 2.3.2.1 where an increase in the quantity of dimensions for a product or price from one to two attributes causes a significant reduction in consumer sophistication but the marginal effect on consumer sophistication of using three or four attributes is smaller (Kalayci and Potters, 2011).

Cumulatively, these three articles identify that complexity can reduce the optimality of consumer decision making, as hypothesised throughout the obfuscation literature. However, complexity also reduces the probability that an alternative will be chosen when there are less complex options available. Lotteries also fit particularly well with many decisions that consumers are expected to make. For example, when purchasing car insurance, an obvious example of Spiegler's (2006) framework, individuals must evaluate their subjective probability of an accident and its likely cause, the probability of theft, the likelihood of windscreen damage, the expected annual mileage, their valuation of the vehicle, and a range of other attributes. Consumers must also trade-off increases in the initial premium with the policy excess. This fits well with a lottery in which future outcomes are uncertain but individuals have some idea of the possible outcomes and their probabilities.

Sitzia & Zizzo (2011) develop two experiments that use lotteries to elicit individuals' preferences towards complex products. One main divergence, however, is that in some of the treatments the price and quantity available are chosen by a human seller before being revealed to the human buyers. Buyers must also decide how much to buy, rather than a binary decision of whether to buy. Sitzia & Zizzo (2011) use a binary measure of complexity. Some products are labelled simple, whilst others are labelled complex. The distinction between the two *products* is that simple lotteries have a small number of possible outcomes and complex lotteries are compound lotteries derived from the simple lotteries.

Each lottery has three possible outcomes, where p_i and x_i are the probability and payoff of alternative i . Therefore, the simple version of lottery j is: $L_{sj} = (p_1, x_1; p_2, x_2; p_3, x_3)$. The complex lottery assigns weights to the outcomes of three independent draws of the lottery L_{sj} such that: $L_{cj} = \alpha L_{sj} + \beta L_{sj} + (1 - \alpha - \beta)L_{sj}$. Therefore, the simple lottery has only 3 potential outcomes, whereas the complex lottery has up to 27.

In the first experiment, three treatments are simulated and in all treatments some subjects play the role of buyers whilst others play the role of sellers. The only difference across the treatments is the sophistication of the seller, which is controlled by changing the amount of information and practice that sellers receive. In the baseline treatment, one seller chooses the price and available

quantity for four buyers. There is no bidding amongst buyers and the price is fixed. Demand is also fulfilled on a first come first served basis such that some buyers may not have the option to purchase if other buyers choose to purchase all of the stock. The amount of available stock is not known to buyers. Whilst this is not the focus of their study, imposing stock constraints introduces competition between the buyers as they must buy quickly before the stock is depleted.

The objective of each seller is to maximise profit, defined as revenue minus the marginal cost of supplying the product. Costs are known to sellers throughout. Sellers do not choose the lottery frame in which prices are presented. Instead, some buyers receive the simple lottery for the first two periods, whilst others receive the complex lottery. This reverses in the second two periods. Buyers are given a float of 390 points to spend in each period. Credit is not transferable across periods but any unused points are counted at the end of the experiment and transferred into cash. The lotteries are also drawn at the end of the experiment and buyers receive a scaled amount of what they win.

The second and third treatments involve sellers that are more informed (IS). In treatment two (IS1), sellers are given more information to develop their pricing strategy. They also play six periods as a buyer and are advised of the existence of complexity aversion and other complexity-related behaviours of consumers. IS2 is identical to IS1 except both the simple and complex lotteries are available to buyers simultaneously, rather than only one simple or one complex option being available (Sitzia and Zizzo, 2011).

Experiment 2 is a test of individual choice (IC) with a computerised seller that draws prices randomly. Prices for some buyers are drawn from the interval [75, 95], whilst prices for others are drawn from [45, 65]. Buyers are aware that prices are computerised and random but they do not know the distribution, which is uniform. Quantity is also infinite to remove the rationing effects in experiment 1. 268 individuals participate across the five treatments.

They find that average prices are marginally higher for complex products in experiment 1, which supports the idea of Carlin (2009): Higher prices will be combined with higher complexity. However, the explanation is different here because there is no competition amongst sellers. They also find evidence of complexity aversion in the short run. Prices in experiment 1 decrease over time for the simple product but remain above the seller's cost, which is expected as there are no competitors. Elasticities also show that individuals are more sensitive to price changes for complex products and their explanation is that individuals pay closer attention to prices of complex lotteries. Treatment IC2, in which two products are simultaneously available, is the only part of the experiment that reveals a statistically significant effect of complexity exploitation in the sense that buyers' purchasing quantities are distorted. Demand is 15% lower for simple products than complex products, such that complexity induces consumers to buy more.

To develop these insights they conduct a series of random effects regressions, similar to Sonino et al. (2002). They find that demand declines as the price increases and demand is higher in the IC treatments. Surprisingly, the complexity parameter is not significant in many of the regressions but the interaction term between complexity and an indicator for the IC treatment is significant and positive. Therefore, complexity distorts purchasing quantities more than it distorts the decision of whether to buy.

Combining high prices with high complexity generates higher sales for the seller than if prices were combined with low complexity. This reinforces the idea that complexity causes several inter-related effects on consumer choice and market competition. One limitation of this study is the absence of competition amongst sellers. Buyers choose only how much to buy, not who to buy from. It is, therefore, not clear that these results will be directly transferable into competitive markets and raises the question of whether complexity aversion will strengthen as competition is introduced?

Iyengar and Kamenica (2010) provide evidence that complexity aversion grows when consumers are given more choice. They continue to model complexity as the ease with which individuals can calculate the expected outcome of lotteries but their focus is different. Instead of examining the responses of individuals to simulated variations in complexity, they study the probability that individuals select the simple alternative as the amount of risky alternatives increases. In addition, they study a combination of both laboratory and field experiments that analyse the allocation of retirement investment decisions between simple and complex funds, as the number of available funds varies.

The first part of their study consists of two experiments. Individuals were approached by research assistants in the vicinity of Columbia University campus and asked to complete a short questionnaire, unrelated to the main experiment, before completing one of two possible experiments. In experiment 1, participants select one gamble and the outcome of the gamble is their compensation for participating. Some participants choose from 11 gambles, others were offered 3. One gamble is riskless whilst the remainder are risky. For the risky gambles, the payoff is conditional on the binary outcome of a coin toss. Therefore, individuals must choose one of the gambles before tossing a coin to determine their winnings. The riskless gamble is presented in the same framing as the risky gambles, except the payoff for both possible faces of the coin are the same. Notably, the option with the lowest risk is also the least complex.

They find that just 16% of 69 individuals choose the simple lottery when three alternatives are offered but 63% of 68 individuals choose the simple lottery when 11 alternatives are offered, demonstrating a significant increase in complexity aversion as the choice set expands. However, it is not clear whether this result is driven by risk aversion or simplicity seeking behaviour, which is disentangled in experiment 2 where the most simple option is also most risky.

The payoffs of the eleven gambles in experiment 2 are determined by the roll of a die, such that the individual must select one of the gambles before the die is rolled and the face that shows determines the payoff of each gamble. For ten of the alternatives, each face corresponds to a unique numerical payoff. That is, each of the faces yield different potential payoffs. For the simple gamble, three of the faces yield one payoff, whilst the other three faces yield another payoff. Their results confirm that individuals migrate towards simple options as the choice set enlarges, which cannot be attributed to risk aversion. Specifically, 16% of 62 participants choose the simple alternative in the three alternative treatment but 57% of 58 participants choose the simple option in the eleven alternative treatment.

Iyengar & Kamenica (2010) also conduct a field experiment on the retirement investment decisions of 588, 296 individuals in 2001. Retirement decisions lend well to choice analysis because participation is high, plan administrators cannot guide individuals' decisions and defaults are not available.²⁴ They analyse how individuals allocate their savings across different investment categories, including money market funds, bonds and equities, when the number of funds offered changes. Only 10% of plans offer less than six choices or more than 20 choices, with considerable variation across plans. Equities are more complex than bonds because bonds pay a fixed rate of return. Therefore, the proportion of funds that are allocated to the different components is a yardstick of individuals' preference for simplicity. The data also includes personal information on investors, including gender, age, tenure, income and wealth. Individuals with income outside the range of \$10,000 - \$1,000,000 and individuals who did not pay into their plan were excluded.

Iyengar & Kamenica (2010) highlight a negative relationship between the number of funds available on a plan and the proportion of investments that are assigned to equities, with some of the reduction in equities directed towards bonds. Increasing the number of alternatives by 10 causes the share of investment in bonds to increase by nearly two percentage points, relative to the 6% mean (Iyengar and Kamenica, 2010). On average, 53% of individuals choose not to invest in equities but increasing the quantity of available funds by 10 causes the probability of not investing in equities to increase by 27%. The reduction in the share of investments assigned to equities is especially interesting because equities tend to represent a larger proportion of the available funds in plans that contain more funds. They also find evidence of a decreasing marginal effect of increasing the number of alternatives, on the investment share of equities.

Iyengar and Kamenica (2010) also show that these portfolio adjustments are not motivated by a stronger preference for prominent funds as the degree of choice increases. Survey data indicates that retirement investors are more familiar with company stock funds (equities) than bonds. This is particularly relevant for retail markets because individuals tend to face a wide variety of

²⁴See Agnew and Szykman (2005) for a laboratory experiment in which the proportion of retirement investments allocated to a default fund may increase with the complexity of the decision, depending on the sophistication of the individual.

choice for seemingly homogeneous products, rather than one or two different options. Therefore, experiments that include only one or two alternatives, including Sitzia and Zizzo (2011), may underestimate individuals' aversion to complexity.

The experiment most closely aligned with retail markets is the recent work of Homburg et al. (2014). They study how changing the complexity of a price scheme influences consumers' estimates of the prices. They also investigate the origins of consumers' behavioural reactions and aversion to complex prices, which is an under-researched topic. They address two key questions. Firstly, are simple prices perceived to be lower prices and more favourable to consumers? Secondly, why do consumers perceive simple prices to be more attractive than complex prices? Their measure of complexity also extends beyond multi-dimensional pricing structures that are at play in lottery-based studies. They use the number of dimensions, the variety of numbers and the amount of different calculations required to compute the total price, as barometers of complexity.

They report two experiments. In the first, they analyse how changes in the objective complexity of a price plan influences consumers' perceptions of the complexity of the price, the perceived price fairness, the perceived transparency of the pricing strategy and their willingness to buy. The idea of price fairness is deeply rooted in marketing philosophy and captures consumers' perceptions of whether the price appears 'right, just and legitimate' (Campbell, 2007; Homburg et al., 2014; Xia et al., 2004). This differs from the traditional competition law interpretation of unfair prices as exploitative prices that become too high with respect to a reference price, when the state of the market changes (Akman and Garrod, 2011). This might include a change in the cost of supply or a reduction in competition, which allows firms to achieve an excessive price-cost margin. Alternatively, prices can be categorised as unfair if they are excessively low, if suppliers seek to erode long-term competition using predatory pricing. Therefore, in law and economics, unfair pricing tends to refer to the price level rather than its framing.²⁵

The subjective transparency of prices in Homburg et al (2014) also refers to whether the consumer perceives the pricing to be "intentionally open, understandable and honest" (Homburg et al., 2014, p. 1116). Encompassing these elements is critical because consumers' perceptions of the pricing scheme will influence choice. Homburg et al. (2014) hypothesise that increasing the complexity of a price scheme will reduce the perceived transparency of the firm and the perceived fairness of the contract. They also contend that perceived transparency and perceived fairness will be positively correlated with consumers' willingness to buy.

²⁵See Akman and Garrod (2011) for a detailed review of the topic of unfair pricing in law and economics and possible methods for identifying unfair price levels. In this field, "prices should only be found to be unfair if they are unfair due to competition issues. Consequently, an abuse should only be found if the firm gains sufficiently at the expense of its customers due to a lack of competition" (Akman and Garrod, 2011, p.418). See Xia et al. (2004) for a discussion of unfair pricing from a marketing discipline perspective. They discuss how prices that are perceived to be unfair can create negative emotions of anger and dissatisfaction, which translate into aversion towards purchasing that alternative.

In the first experiment, 385 undergraduate students are presented with one of eight different mobile phone tariffs and asked to evaluate the tariff according to a fictional usage profile. The tariffs vary in their objective complexity and each participant receives only one tariff. Participants are remunerated by entry into a fair lottery for a book voucher, which is independent of their responses during the experiment. Individuals report their evaluation of the complexity and fairness of the tariff and their perception of the transparency of the seller's pricing strategy, which is measured on a four-point scale. Information for a series of control variables and details of the participant's actual mobile phone usage is also recorded.

They find that individuals' perceptions of tariff complexity increases with each of the three attributes used to simulate tariff complexity. Using a structural equation model, they confirm that price complexity causally reduces individuals' perceptions of the transparency of the pricing strategy used by the seller, which reduces perceptions of price fairness. More importantly, individuals report a lower willingness to buy the contract when the perceived transparency of the seller and the perceived fairness of the pricing structure are lower. Homburg et al. (2014) also find no evidence that consumers view simple tariffs as suspicious, which might generate an aversion to the most simple tariffs.

Their second experiment is more closely aligned with Huck and Weizsäcker (1999) and Sonino et al. (2002), as participants state their relative preferences for two tariffs that are presented simultaneously. Note that the subject group is different for the second experiment and includes 176 individuals with an approximately equal share of students and non-students. In stage 1, participants are presented with one simple and one complex mobile phone tariff and asked to report the perceived complexity and price fairness associated with each tariff, and the perceived transparency of each pricing strategy. In stage 2, participants estimate the total price for the different tariffs, conditional on the same fictional consumption patterns as experiment 1. In stage 3, participants report information for similar control variables as in the first experiment.

The price of each tariff is randomly set at either €30 or €35 such that there are three possible relative price combinations: The simple product is cheaper, the complex product is cheaper or both tariffs have the same price. A difference-in-difference approach is then applied to calculate the difference between the participants estimates of the differences between the tariffs and the actual differences in prices. This directly tests for any systematic bias in the estimation procedure due to complexity.

Descriptive statistics indicate that the objectively less complex tariff is perceived to be the simple option. More importantly, individuals perceive simple prices to be more fair and associated with a more transparent underlying pricing strategy, independently of the actual price of the contract. Therefore, tariffs that are less complex but more expensive, are perceived to be more fair and transparent. In addition, when the simple tariff is more expensive by €5, the estimated price

difference is only a quarter of the actual difference, at €1.26. Therefore, consumers consistently estimate that simple prices are cheaper.

A structural equation model also reinforces the causal relationships identified in the first experiment. They document a causal association in which an increase in the perceived complexity of the tariff reduces the perceived transparency of the pricing strategy, which reduces the perceived price fairness. Increases in the perceived transparency and price fairness of a tariff also increase the probability that the tariff will be chosen. Increasing the perceived price complexity also reduces the magnitude of the estimated price advantage by a statistically significant amount. In particular, 65.9% of individuals choose the simple tariff when it is €5 more expensive. This suggests that consumers who cannot accurately estimate the prices are systematically biased towards the least complex tariff available, rather than purchasing randomly. These results cast doubt on the validity of competitive pricing models grounded on a random purchase rule.

2.5 Conclusion

In this review I have delivered a comprehensive analysis of the theory and evidence relating to the intersection between price competition, obfuscation and consumer behaviour. Several core themes have emerged. Firstly, scarcity plays a key role in consumer choice. Consumers are limited in their resources or inclination to engage with firms, which sellers can subsequently exploit or manipulate to sustain supra-competitive profits. Secondly, obfuscation creates market inefficiencies and these costs are often borne by society. Thirdly, obfuscation takes a wide variety of forms and is intimately linked to topics of salience, price-framing, advertising and price updating. Fourthly, policies to limit the impact of obfuscation, either using demand-side initiatives such as price comparison aids or education programmes, or supply side policies to prohibit obfuscation practices, may not always generate the intended social gains. This motivates a thorough understanding of firms' incentives before such interventions can be considered. Fifthly, the assumption that consumers who cannot compare prices buy randomly is a poor approximation for real world behaviour and future research should readdress this concern, which is a central objective of this thesis.

More broadly, obfuscation as a research area emphasises that the structure of prices is more important than a straight-forward marketing or price promotion decision for firms. Instead, the structure and presentation of prices can affect earnings significantly. Several additional strands of literature are clearly related to the topics of this review. Firstly, I have not considered the emerging literature in which individuals choose to simplify their choice procedure, rather than selecting less complex alternatives. In those studies, individuals choose to simplify the choice process itself before applying well-defined preferences to the remaining choice set. Bachi and Spiegel (2015) and Papi (2014) present promising research in this area.

This review also does not consider the impact of the effect of policy initiatives in significant detail. The reason for delaying a direct examination of the policy literature is threefold. Firstly, in many cases the implications for policy are clear. For example, in Zhou's (2014) model of multi-product retailing where search costs benefit consumers, policies that reduce the search cost can intuitively harm them. Secondly, policies tend to be specific to individual market problems. For example, a policy designed to increase consumer sophistication using education initiatives does not address the issue that firms can change their prices excessively frequently. Thirdly, the effect of remedies to obfuscation and manipulative pricing are considered in significant detail in the following two articles contained in this thesis.

The literature on advertising and, in particular, false advertising is also relevant to this research area because firms may impair consumer decision-making by presenting factually incorrect materials. Advertising is a broad area of research within Industrial Organisation and is an ideal candi-

date for a future review. Several applications of false advertising are discussed in the subsequent chapters of this thesis.

Chapter 3

Competition with Simplicity Biased Consumers

3.1 Introduction

In markets where firms supply identical products, consumers should have little difficulty in identifying the best deal. When this occurs, competition should discipline firms to price at marginal cost. In practice, however, prices persistently exceed the competitive level, which in the absence of product differentiation is often attributed to imperfect consumer decision-making. This could arise from cognitive limitations or inattention. A growing concern for regulators is that firms inject spurious complexity into their products to impair consumer decision-making further and exacerbate disengagement, referred to as strategic obfuscation. Examples include excessively fragmented multi-part tariffs and complicated framing of product information.

This article explores a potential caveat of obfuscation for firms, which makes such practices unappealing without requiring regulation. The crux is that consumers prefer simple options, particularly when they become confused. Therefore, when a firm injects complexity into its offer it may inadvertently make its product less appealing to the now larger group of consumers who are unable to comprehend prices. We examine how this preference for simple options by consumers influences the way in which firms price and promote their products. To the best of our knowledge this is the first article to study the interaction between simplicity biased consumers and competitive obfuscation.¹

Specifically, we address the following important questions: How does the presence of simplicity biased consumers affect prices, profits and consumer welfare? Does the presence of simplicity biased consumers enhance or detract from market transparency? Do policies to remove behavioural biases and improve consumers' resilience to complexity improve consumer welfare?

¹Chioveanu and Zhou (2013) analyse an interesting model which features an opportunity for consumers to prefer products that are presented in a simple way but their focus is different. This is discussed in the literature review.

What is the correlation between prices and complexity? When should consumers select the simplest option?

To answer these questions we analyse a three period price complexity model in which firms choose the complexity of their product before competing in prices. Consumers subsequently make their purchasing decisions where there are two kinds of consumers: Informed and Uninformed. Informed consumers have perfect knowledge of prices and purchase from the lowest priced seller. Uninformed consumers are impeded by market complexity and cannot compare prices.

The proportion of informed consumers is determined endogenously as a decreasing function of firms' obfuscation efforts such that when the market is more complicated fewer consumers are able to understand prices. The main departure from the existing literature lies in the choice procedure of uninformed consumers, who are often assumed to buy randomly. Instead, this model allows consumers who cannot compare prices to be biased towards simple products, captured by some uninformed consumers purchasing from the least complex seller.

The obfuscation incentives of firms can be summarised as follows. Each firm wishes to obfuscate to stimulate the proportion of uninformed consumers, which softens price competition. Simultaneously, each firm wishes to offer a product less complex than their rival to attract simplicity biased consumers. Therefore when consumers are assumed to be indifferent between firms of different complexity, all firms select maximum complexity before competing in prices.

The following example captures the intuition. Suppose there are five consumers choosing a mobile phone contract and market complexity confuses four of them. The informed buyer will identify and purchase the cheapest tariff but uninformed consumers cannot determine which firm offers the lowest price. In this situation the random purchase rule, prevalent in the literature, implies that uninformed consumers buy randomly but in our model we recognise that uninformed consumers are biased towards the simplest tariff. Therefore some uninformed consumers buy from the least complex firm and the remainder buy at random.²

The introduction of simplicity biased consumers is motivated by recent experimental and empirical studies that report a systematic bias by consumers towards simple offers (McGovern and Moon, 2007). Most closely related is Homburg et al. (2014) who report two experimental studies regarding the impact of price complexity on consumers' purchasing decisions. They use the quantity of elements in a price plan, the variety of numbers and the variety of calculations required to calculate the total price as proxies for complexity. In the first experiment they find that more complex pricing leads a firm to be perceived as less transparent with lower price fairness. In the second they show that buyers systematically overestimate the price of more complex tariffs

²In the model this is captured by a constant proportional rule such that a constant fraction of uninformed consumers buy from the least complex seller.

and underestimate the price of less complex tariffs. This leads consumers to consistently purchase less complex but more expensive tariffs.³

Iyengar and Kamenica (2010) also investigate experimentally and empirically how retirement investment decisions are influenced by the degree of choice available, which may be interpreted as decision complexity. They find that individuals prefer simple investments and this preference strengthens as the decision becomes more taxing due to an increase in choice. Saliently the preference for a simple financial product is characterised by consumers selecting products that are intrinsically simple, rather than funds managed by prominent brands. This distinguishes simplicity biases from brand preferences. De los Santos et al. (2012) also find evidence that non-price factors, such as the complexity of an online interface and ease of a transaction, strongly influence consumers' preferences in an empirical analysis of online book sellers.

There are also theoretical reasons for consumers to prioritise simple products because many theoretical and experimental studies predict that obfuscating firms charge higher prices when buyers engage in 'all-or-nothing' search (Carlin, 2009; Gu and Wenzel, 2014).⁴ It is reasonable to assume that upon repeated interaction, consumers who cannot observe prices may infer the relative price from a firm's obfuscation.⁵ From a different perspective, Gaudeul and Sugden (2012) show that consumers are best served by confining their consideration to firms that present their product in accordance with an industry standard. Firms that deviate to individualised promotion formats, which are more complicated for the consumer to interpret, charge higher prices in equilibrium. Buyers may also prefer transparent firms when they are concerned with price fairness.⁶

Consumers' preferences for simplicity are also evident from firms' promotion strategies. Sainsbury's, a leading UK supermarket, recently advertised a joint-venture with Telecoms provider Vodafone to create 'Mobile by Sainsbury's.' In an industry with minimal product differentiation, Sainsbury's champion simplicity for consumers as their principal selling point by removing excessive choice and unnecessary jargon.⁷ Price comparison websites also champion simplicity for the consumer as a feature of both their service and their brand. For example, *Comparethemarket.com*, a leading UK price comparison website for energy, finance and insurances, claim; "We're here to make saving money and getting a better deal as simple as possible."

³Within the marketing, economics and psychology literatures there are many explanations for the simplicity bias. However, the model presented in this study is general and not predicated on any particular source of simplicity bias. See Homburg et al. (2014) and the references contained within for further discussion on the role of marketing in consumer decision-making.

⁴See Kalayci (2015) for an experimental examination of Carlin (2009) and Gu and Wenzel (2015) for an experiment in which complexity and pricing are positively correlated when firms vary in their prominence.

⁵In the present model this intuition fails because firms that offer simple products charge higher prices as a direct consequence of the simplicity bias.

⁶See Homburg et al. (2014) and the references within.

⁷See www.telegraph.co.uk/finance/newsbysector/mediatechnologyandtelecoms/10152166/Sainsburys-vows-change-as-it-launches-mobile-phone-service.html. Accessed 04/05/16.

Obfuscation in this article differs strongly from traditional product differentiation. In that literature, firms decide whether to invest in a higher quality product to attract quality-sensitive consumers, or produce a lower quality product, at lower cost, and sell only to quality-insensitive consumers. Both obfuscation and product differentiation influence the allocation of demand because complexity determines how uninformed consumers purchase. However, obfuscation induces a secondary effect by influencing the composition of the consumer population in terms of informed and uninformed buyers.

3.1.1 Summary of Main Results

We show that in the absence of simplicity biased consumers, firms universally select high complexity. There exists no incentive to compete in obfuscation and industry profit, to which each firm expects to receive an equal share, is strictly increasing in the degree of consumer confusion. By comparison, when the preference for simplicity is introduced this equilibrium is overturned under very general conditions. Firms now seek to undercut the obfuscation level of their rival to attract simplicity biased consumers, which produces a mixed strategy equilibrium. This leads to greater market transparency but the degree of transparency depends on three factors.

First, the strength of the simplicity bias. Second, the natural level of informed consumers when firms do not obfuscate. Third, the rate at which obfuscation impairs consumers' ability to compare prices. Transparency is maximised when the simplicity bias is large, the natural level of informed consumers is low and the effect of obfuscation on consumer sophistication is small. For a sufficiently strong bias and sufficiently small marginal effect of complexity on consumer sophistication, all firms are disciplined to select low complexity. This is interesting because the simplicity bias appears inactive: all uninformed buyers select randomly, but firms find obfuscation prohibitively costly.

Most strikingly the model demonstrates that when consumers who cannot identify the lowest priced seller are biased towards a simple offer they will pay the highest price in the market but surprisingly firms may earn lower profit as the degree of bias increases. This paradox arises because the presence of simplicity biased consumers incentivises firms to compete directly in their obfuscation, which leads to greater market transparency. As a result more consumers are able to find the cheapest seller and price competition between firms intensifies. However, in equilibrium the firm with the lowest complexity charges the highest price because the least complex firm is guaranteed a larger share of uninformed consumers, who purchase at any price not exceeding their reservation threshold. Therefore, any reduction in price to compete for informed buyers is more costly for the least complex firm.

In every case firms select prices stochastically according to a probability distribution to prevent the usual marginal undercutting deviation by a rival, which provides an explanation for the preva-

lence of firms changing both their complexity and price over time. These results are particularly interesting for competition and consumer protection authorities as policies to reduce consumers' behavioural biases will generate higher obfuscation by firms and may increase prices. The results also highlight the importance of considering how consumers may respond to the behaviour of firms.

This article also illuminates a novel role for false advertising if firms are able to tamper with consumers' perception of complexity.⁸ By increasing the complexity of a product but using strategic marketing that persuades buyers that their offer is simple, firms can simultaneously increase the proportion of uninformed buyers and attract simplicity biased consumers. This introduces an additional dimension to individualized promotion strategies as firms must now appear simple to confused consumers, in addition to adopting a complicated promotion format.

For example, Telecoms operator O_2 offer a 'simplicity' tariff that divorces the cost of a mobile handset from monthly usage charges. This product can be positively explained as a response to heterogeneous consumer preferences, which naturally leads to product differentiation. However, the tariff structure is less comparable with rivals and marketed as a less complex contract for consumers. In the model we also describe how this intuition might extend to multi-brand retailing where firms stimulate market complexity using one brand and lure simplicity biased consumers under another. This differs from conventional multi-outlet retailing because firms are typically predicted to adopt the same strategy across outlets (Ireland, 2007).

3.1.2 Related Literature

This article contributes to the growing literature on strategic obfuscation in which firms inject spurious complexity into retail markets and conceal information from buyers. Ellison (2005) shows that a firm may lure consumers using a low quality product before encouraging them to upgrade to a higher quality product that is less competitively priced.⁹ Gabaix and Laibson (2006) show that a firm may shroud the existence of add-on products and their high prices until the consumer has purchased the base good. Spiegler (2006) also shows that a firm may increase the quantity of dimensions for a single product and price each element with large variation, when consumers base their decisions on only a randomly chosen subset of the products attributes.¹⁰

The most important distinction with this literature is that in these models firms have no incentive to undercut the complexity of their rival because consumers are indifferent between obfuscating and non-obfuscating firms. In contrast, this article introduces an incentive for firms to obfuscate less than a rival to attract simplicity biased consumers. In addition, obfuscation in the

⁸First steps are made in this direction in Chapter 3 and included more formally in Chapter 4.

⁹See Ellison and Ellison (2009) for an empirical application to the market for computing memory.

¹⁰For example, home insurance contains many dimensions and a consumer cannot reasonably evaluate every policy on every dimension. Therefore, consumers use a particular dimension when appraising competitors' prices.

present model is not predicated on the existence of add-on items, upgrades or multi-dimensional products. Rather, obfuscation is interpreted most broadly as any form of unnecessary complexity a firm may introduce into its product to confuse consumers.

Carlin (2009) presents a model in which firms add complexity to their prices to reduce the proportion of informed consumers (experts). The model shows that firms become increasingly reliant on complexity when competition intensifies. In Carlin (2009), however, uninformed consumers also follow a random purchasing rule with no bias towards firms that simplify their pricing.

Interestingly, Dahremoller (2013) shows that when firms vary in their add-on profitability the incentive for all firms to obfuscate is overturned in Gabaix and Laibson (2006), providing the decision to shroud precedes price setting. The timing is similar to the present study and in both frameworks obfuscation may not be optimal but the intuition is different. In Dahremoller (2013) the efficient or inefficient firm may unshroud depending on the underlying demand structure and specifically the proportion of uninformed consumers when firms do not obfuscate. In our study firms are symmetric and sell only a single product but the shrouding equilibrium is overturned because consumers favour simple products.

The model also relates to the established search cost literature (Ellison and Wolitzky, 2012; Wilson, 2010). Wilson (2010) shows that in an environment with costless and costly search consumers a duopolist may profitably increase the time required for consumers to evaluate its product. The intuition follows that by increasing the search cost, costly search consumers are directed to the more prominent rival, who subsequently charges a higher price due to the reduction in competitive pressure. The obfuscating firm charges a lower price to secure the custom of costless search consumers but interestingly profit for both firms exceeds the no-obfuscation level. The model we study features some similar dynamics however the intuition is different. Instead of a firm increasing its obfuscation to soften competition on the rival, firms now start from a position of maximum obfuscation, if all uninformed consumers buy randomly. Firms subsequently reduce their obfuscation symmetrically in response to an increase in the degree of simplicity bias and surprisingly consumer welfare may increase with the degree of bias depending on the level of consumer sophistication.

Gu and Wenzel (2014) present a related model in which firms also select a level of obfuscation before competing in prices, and obfuscation determines the proportion of informed and uninformed consumers. However, in their model firms vary in prominence. This is captured by one firm receiving a larger share of uninformed buyers independently of its obfuscation and price. In our model firms are *ex ante* symmetric but the distribution of uninformed consumers may become asymmetric at the pricing stage if firms select different complexity levels because uninformed consumers are biased towards the least complex seller. In this sense the relative obfuscation of each firm determines their prominence endogenously in this framework, which links to Wilson (2010).

These differences are also reflected in the correlation between obfuscation and prices. In their model the prominent firm selects maximum obfuscation and charges the highest average price but in ours the firm with the lowest obfuscation charges the highest price because consumers are simplicity biased.

Gu and Wenzel (2014) also highlight that policies to limit the effectiveness of obfuscation on consumer sophistication may be ineffective as the less prominent firm may increase obfuscation to counteract the policy. We reinforce this concern in a symmetric setting as policies to increase the proportion of informed consumers may lead to increased obfuscation by both firms, depending on parameters within the model.

The model also develops our understanding of price-frame competition (Eliaz and Spiegler, 2011a; Piccione and Spiegler, 2012). Most relevant is Chioveanu & Zhou (2013, henceforth CZ) who develop a comprehensive framework to explore the role of different sources of consumer confusion on firm behaviour. Using a duopoly model CZ show that when consumers are impaired by complex framing and firms choose a complexity level before setting the price, all firms adopt maximum complexity. The present model captures this equilibrium as the special case when consumers do not prefer less complex options. In contrast, the introduction of simplicity biased consumers in this article overturns this solution as firms compete in complexity.

CZ also develop their model to consider an oligopoly environment and alternative decision procedures of consumers, including a preference for firms that adopt a less complex frame. However, we confine our focus to a duopoly environment to explore the role of simplicity biases in greater detail, the impact on firms' incentives and the surprising results for consumer welfare. Our model also identifies additional opportunities for firms to manipulate simplicity biased decision-makers, which are particularly acute for competition and consumer protection authorities.

This article also relates to the literature on the role of common standards in retail markets.¹¹ Gaudeul and Sugden (2012) show that firms may choose to conform to industry standards that enhance consumer decision-making when consumers discount individuated promotions. They also find that when a fraction of firms follow a common standard and the remainder adopt an individuated format, consumers are rational in discounting firms that adopt individuated formats as these firms charge higher prices.

In our framework and Gaudeul and Sugden (2012), demand-side pressures improve market transparency and can increase consumer welfare but the intuitions differ strongly. In theirs, the heuristic based decision rule is compatible with utility maximisation in the sense that firms conforming to a common standard charge lower prices. However, in the present model consumers who select the simplest option pay the highest price, which is incompatible with rational utility

¹¹For example, sugar and milk are complimentary items in UK coffee stores but syrups and cream are chargeable extras.

maximisation if utility is solely determined by prices. In Section 3.6 we discuss several explanations for the persistence of simplicity-seeking behaviour by consumers when more complex options are cheaper based on, amongst other factors, imperfect consumer decision-making and additional non-monetary utility derived by consumers that obtain an intrinsically simple product. Therefore, this article aids our understanding of why consumers may consistently purchase more expensive but less complex products (Homburg et al., 2014).

The model also provides a novel explanation for firms stochastically changing their promotions over time. Carlin and Manso (2011) show that when more consumers become informed over time due to learning effects, firms will strategically refresh their menu, at cost, to restore the original level of consumer sophistication and soften price competition. In our model firms change their promotions over time but the intuition is very different as firms now change both their menu and the complexity of their menu to compete for simplicity biased consumers.

Broadly, this article contributes to the economic interest in choice architecture. Bachi and Spiegler (2015) and Papi (2014) consider decision-making when individuals exhibit tradeoff avoidance. They show that firms earn profits in otherwise competitive markets because decision makers are forced to neglect important attributes to simplify their decision-making. The distinction is that in these models decision-makers simplify their decision-making procedure to evaluate competing offers whereas in the present model the simplicity bias forms a key part of the choice procedure as ‘all-or-nothing’ type buyers are biased towards simple options. Finally, this study contributes to the strengthening literature on behavioural industrial organization and bounded rationality, which explores the competitive implications of non-standard approaches to consumer decision-making.¹²

The article is organised as follows. Section 3.2 outlines the main model. Equilibrium pricing is characterised in Section 3.3 and obfuscation strategies are determined in Section 3.4. Consumer welfare and industry profit is analysed in Section 3.5. In Section 3.6 we discuss the main results, consider the implications for consumers, firms and policy and highlight testable predictions for empirical study. Section 3.7 concludes.

3.2 A Model of Competition with Simplicity Biased Consumers

Consider a game of non-cooperative competition with perfect information between two firms indexed $i, j \in N$, who supply a homogeneous product to a unit mass population of consumers. The firms produce at a constant symmetric marginal cost which is normalised to zero and there are no capacity constraints on the output of either firm. The sellers are also equally prominent and consumers have no pre-existing brand preferences.

¹²See Spiegler (2011) for a textbook treatment.

Each consumer will purchase one unit of the product providing the price does not exceed their common reservation threshold, $r = 1$. At stage 1 firms simultaneously and independently select low or high complexity (k), which determines the complexity of their product.¹³ In stage 2 the firms must select the price of their product and in stage 3 consumers purchase according to their information type and bias.

Consumers vary only in their ability to evaluate the firms' offer when it is presented in a complicated way. Moreover for a given level of market complexity each consumer is either informed or uninformed, where the proportion of informed consumers is a decreasing function of the total complexity created by firms. Informed consumers, $\mu_{(k_i, k_j)} \in (0, 1)$, compare prices and buy from the lowest priced supplier. Uninformed consumers ($1 - \mu_{(k_i, k_j)}$) cannot compare anything and within the literature they typically purchase at random. In our model, some of the uninformed consumers will instead purchase the simplest option with a symmetric tie-breaking rule. Specifically, fraction $\delta(1 - \mu_{(k_i, k_j)})$ purchase the least complex option whilst $(1 - \delta)(1 - \mu_{(k_i, k_j)})$ purchase randomly, where $\delta \in [0, 1]$ is constant, exogenous and common knowledge throughout the game.

Uninformed consumers are not price-sensitive but will never pay more than their common reservation price. This can be interpreted as an ex-post sales constraint whereby consumers could return a product that is realised to be priced above the reservation threshold. For any level of obfuscation there exists a positive proportion of uninformed consumers, which is consistent with empirical evidence (Wilson and Waddams-Price, 2010). The game is summarised by the following timeline:

Period 1:

- Each firm $i \in N$ selects (without cost) a level of complexity, $k_i \in \{L, H\}$, with which to present their product.
- The proportion of informed consumers is determined as a declining function of aggregate complexity.
- The relative complexity of each firm becomes common knowledge and determines the allocation of uninformed consumers.

Period 2:

- Each firm $i \in N$ selects a price (distribution).
- The relative price of each firm determines the allocation of informed consumers.

¹³For generality the complexity of a product refers to any spurious complexity a firm may introduce into the product or the presentation of its price.

Period 3:

- Consumers purchase according to their type:
- $\mu_{(k_i, k_j)} \in (0, 1)$ buy from the lowest priced firm.
- $\delta(1 - \mu_{(k_i, k_j)})$, $\delta \in [0, 1]$, buy from the least complex firm.
- $(1 - \delta)(1 - \mu_{(k_i, k_j)})$ buy from a randomly chosen firm.

The timing of firms selecting the complexity of their product before selecting the price captures the idea that the presentation of a product and its attributes requires a greater time to adjust than the price, which is widely recognised in the literature (Dahremoller, 2013; Gu and Wenzel, 2014, 2015). This fits well with many interpretations of complexity such as the ease with which a consumer can navigate a website, the complexity of a multi-part tariff and the quantity of terms and conditions presented with the product. In addition, the sequential timing makes the impact of simplicity biases most visible. Many of the results continue to hold when prices and complexity are chosen simultaneously, as discussed in Section 3.6 and formalised in Chapter 4.

Write the complexity levels of firms i, j as the pair $\{k_i, k_j\}$. There are four possible outcomes of the first stage: $\{L, L\}, \{L, H\}, \{H, L\}, \{H, H\}$. Let the proportion of informed consumers be denoted μ_L if both firms adopt low complexity, μ_H if both firms adopt high complexity and μ if one firm selects high complexity and the other selects low complexity. Complexity reduces consumer sophistication such that $1 > \mu_L > \mu > \mu_H > 0$. Obfuscation in the model therefore occurs whenever a firm deviates from low complexity, which might be interpreted as the ‘natural’ level of complexity associated with the offer of any firm.

Obfuscation in this model adopts two roles. First, the complexity of each firm contributes to market complexity, which determines the proportion of informed consumers. Second, the relative complexity of each firm determines the allocation of uninformed consumers because uninformed buyers prefer simple options. The game is solved by backward induction.

3.3 Equilibrium Pricing

This section characterises the optimal pricing strategy. Firms produce without cost implying that profits follow the structure $\pi_i = p_i q_i$ for each firm $i \in N$, where q_i corresponds to demand. The firm’s decision is to select a price that maximises expected profit, conditional on market complexity and the relative complexity of each firm. The following definition aids the discussion.

Definition 1: Captive Consumers. The captive consumers for a firm correspond to the proportion of the consumer population that are guaranteed to purchase from that firm once obfuscation decisions have been committed, at any price not exceeding the reservation threshold.

For example, when firms are asymmetric in their complexity, the least complex firm is guaranteed all simplicity biased consumers in addition to an equal share of random purchasers. The more complex rival is guaranteed only the equal share of random purchasers. Alternatively, when firms select the same complexity level the value of captive consumers is equal across firms. The allocation of captive consumers will prove fundamental in determining the optimal price strategy of each firm.

The asymmetric case mirrors Narasimhan (1988) but the intuition is different. In Narasimhan (1988) uninformed consumers are distributed unevenly between the two firms because one firm holds more brand loyal consumers and firms subsequently compete in prices to attract informed consumers. In the present model the distribution of uninformed consumers may become unequal at the pricing stage because one firm is favoured by simplicity biased consumers. In this sense the model incorporates the brand loyalty effect of Narasimhan (1988) endogenously as a function of firms obfuscation decisions in period 1. The symmetric case captures the spirit of Varian (1980) except the proportion of informed consumers is now determined endogenously.

3.3.1 Symmetric Pricing

Consider first the case in which both firms select low complexity. Each firm has secured an equal quantity of captive consumers, which produces identical pricing incentives.

Proposition 1. There exists no pure strategy Nash equilibrium in prices.

All proofs are contained in the appendix. The intuition follows that selecting a price deterministically exposes the firm to the usual marginal undercutting deviation by a rival, which costs the firm any demand from informed consumers.

Proposition 2. For any $\mu_{(k_i, k_j)} \in (0, 1)$ all firms earn positive profits and compete in price distributions bounded from above by r .

At the reservation price the firm is more expensive than the rival with certainty and attracts only an equal share of uninformed consumers. Therefore any maximum price below r yields no increase in demand but a reduction in profit per customer, fixing the upper bound at r .

The payoff at $p_i = r$ is:

$$\pi_i(p_i, p_j | k_i = k_j = L, p_i = r) = \frac{1 - \mu_L}{2} \cdot r$$

By definition of a mixed strategy the firm must be indifferent amongst any price p contained in the support of the equilibrium price distribution. This yields the following indifference condition where $F^L(p)$ is the (symmetric) price distribution of the rival.

$$\frac{1 - \mu_L}{2} \cdot r = \left[\frac{1 - \mu_L}{2} + \mu_L [1 - F^L(p)] \right] p \quad (3.1)$$

The indifference condition makes clear that a firm faces a tradeoff between charging a higher price to exploit captive demand and a lower price in an attempt to secure informed consumers. In equilibrium these tensions must balance yielding a minimum price $\underline{p}_L = \frac{(1 - \mu_L)r}{1 + \mu_L}$ where $F^L(p) = 0$. Solving (3.1) for $F^L(p)$ yields the symmetric equilibrium price distribution:

$$F^L(p) = 1 - \frac{1 - \mu_L}{2\mu_L} \left[\frac{r - p}{p} \right] \quad (3.2)$$

with support $\left[\frac{(1 - \mu_L)r}{1 + \mu_L}, r \right]$. It is easily verified that $F^L(p) > 0$, $F^L(\underline{p}_L) = 0$, $F^L(r) = 1$, satisfying the conditions for a cumulative density function. It is also continuous and atomless by the usual arguments (See Varian (1980)). By symmetry when both firms select high complexity, prices are drawn according to a different distribution function, $F^H(p)$:

$$F^H(p) = 1 - \frac{1 - \mu_H}{2\mu_H} \left[\frac{r - p}{p} \right] \quad (3.3)$$

with support $\left[\frac{(1 - \mu_H)r}{1 + \mu_H}, r \right]$.

Result 1. *When firms are symmetric in their complexity, each firm selects a price by making a random draw from the associated symmetric price distribution $F^k(p)$ with support $[\underline{p}_k, r]$.*

The distributions also indicate that when market complexity is higher, the lowest equilibrium price is higher and prices are less dispersed. The reason is that firms are less willing to make significant price cuts when there are fewer informed consumers because the potential reward is smaller. The average price charged by the firms is also higher when market complexity is higher, illustrating the collective incentive for firms to obfuscate.

3.3.2 Asymmetric Pricing

Consider now the optimal pricing strategy when firms adopt different complexity levels. The allocation of captive consumers is unequal, which produces a single mass point at the reservation price of the distribution, for the least complex firm (Narasimhan, 1988). As a consequence the more complex rival charges prices strictly below the reservation price. Fix the profit of the least complex firm, firm i for convenience, at the reservation price. The firm attracts only captive consumers, which yields a payoff:

$$\pi_i(p_i, p_j | k_i < k_j) = \frac{(1 + \delta)(1 - \mu)}{2} \cdot r$$

As in the symmetric complexity case the mixed strategy solution must make the firm indifferent between any price in the support, producing the following indifference condition:

$$\frac{(1 + \delta)(1 - \mu)}{2} \cdot r = \left[\frac{(1 + \delta)(1 - \mu)}{2} + \mu[1 - F_j(p)] \right] p \quad (3.4)$$

Solving for $F_j(p)$ yields the probability distribution function over prices for the more complex firm:

$$F_j(p) = 1 - \frac{(1 + \delta)(1 - \mu)}{2\mu} \left[\frac{r - p}{p} \right] \quad (3.5)$$

defined over the support $\left[\frac{(1 + \delta)(1 - \mu)r}{1 + \delta + \mu(1 - \delta)}, r \right)$. The lower bound is taken directly from (3.4) where $F_j(p) = 0$.

To compute the price distribution for the least complex firm we are able to write a similar profit indifference condition for the rival at the minimum price, noting that the reservation price is not strictly contained in its support. At the lower bound of the price distribution the firm is guaranteed the demand of all informed consumers and an equal share of random purchase consumers:¹⁴

$$\left[\frac{(1 - \delta)(1 - \mu)}{2} + \mu \right] \cdot \frac{(1 + \delta)(1 - \mu)r}{1 + \delta + \mu(1 - \delta)} = \left[\frac{(1 - \delta)(1 - \mu)}{2} + \mu[1 - F_i(p)] \right] p$$

Solving for $F_i(p)$ yields the optimal price distribution for the less complex firm:

$$F_i(p) = 1 - \frac{1 - \mu}{2\mu p} \left[\frac{(1 + \delta)(1 - \delta + \mu + \delta\mu)r}{1 + \delta + \mu(1 - \delta)} - (1 - \delta)p \right] \quad (3.6)$$

¹⁴Further details of the derivations are provided in the Appendices.

defined over a near-identical support with the inclusion of the reservation price: $\left[\frac{(1+\delta)(1-\mu)r}{1+\delta+\mu(1-\delta)}, r \right]$ and a mass point at r of value: $\frac{2\delta(1-\mu)}{1+\delta+\mu(1-\delta)}$.

Result 2. *When firms are asymmetric in their complexity, each firm selects prices by taking a random draw according to asymmetric probability distribution functions. The firm with the less complex offer has an atom in its price distribution at r . The more complex rival charges prices strictly below r .*

Remark 1. In equilibrium the least complex firm secures a larger payoff.

It can be demonstrated (See Narasimhan (1988)) that the price distribution of the firm that offers the least complex product is higher in terms of first order stochastic dominance. The intuition follows that the firm with more captive consumers is less willing to compete for informed consumers because the marginal cost of a reduction in price, in terms of the foregone guaranteed profit from captives, is larger for this firm than for the rival. The marginal benefit, in terms of the expected value of informed consumers, remains the same. Therefore, prices and obfuscation are negatively correlated at the firm level due to the simplicity bias exhibited by uninformed consumers. This leads informed buyers to purchase predominantly from the more complex seller.

The mass point of the least complex firm also has a neat interpretation as the value of simplicity biased consumers relative to the maximum possible demand the firm is able to attract if it also offers the lowest price. This incorporates the fact that an equal share of random purchase consumers are always guaranteed to the rival. The mass point also constitutes the frequency with which the headline or Recommended Retail Price (RRP) is charged in equilibrium. Prices below the RRP reflect ‘sale’ prices (Spiegler, 2011).





In equilibrium the least complex firm charges both RRP and ‘sale’ prices but the more complex firm randomises continuously over ‘sale-only’ prices. This is consistent with the observation that in many markets firms use sales with different frequency. Some firms continually offer (frequently changing) ‘sale’ prices whilst other firms offer discounts for only a limited period of time.

The value of the mass point is also increasing in the degree of simplicity bias (δ) and decreasing in the intermediate level of consumer sophistication (μ). The reason is that when the bias is larger, more consumers are captives for the least complex firm. This raises the cost of the firm pricing at ‘sale’ in an attempt to attract informed consumers. From the opposite direction, a reduction in the proportion of informed consumers reduces the value of success in price competition and raises the cost of the ‘sale.’

The following figures illustrate these results where $\mu_H = 0.25, \mu = 0.5, \mu_L = 0.75, \delta = 0.5$ and the reservation price is maintained at 1. Figures 3.1 and 3.2 correspond to the probability distribution schedules over prices when both firms adopt low and high complexity respectively.

The lowest price contained in the support of the probability distribution is higher when firms adopt high complexity and there exists no mass point in either distribution.

Figures 3.3 and 3.4 correspond to the probability distributions for the least complex and most complex firm respectively, when firms commit to different levels of complexity. The lower bound is identical but the schedule for the least complex firm contains a mass point at the reservation price of value 0.29. That is, the least complex firm charges the reservation price with 29% probability.

Figure 3.1: Symmetric Pricing: $\{L, L\}$ Figure 3.2: Symmetric Pricing: $\{H, H\}$ Figure 3.3: Asymmetric Pricing: $\{L, H\}$ Figure 3.4: Asymmetric Pricing: $\{L, H\}$

3.4 Equilibrium Obfuscation

It is now possible to characterise the optimal obfuscation strategy in period 1. Firms face competing incentives. Selecting high complexity softens price competition. A higher complexity level also increases the amount of simplicity biased consumers that the firm will receive, if the firm offers a lower complexity level than the rival. Selecting low complexity, however, increases the probability that the firm will secure the demand of simplicity biased consumers. This incentive structure produces a novel tradeoff. The relative strength of these incentives depends on the parameters governing consumer sophistication (μ_L, μ, μ_H) and the magnitude of the simplicity bias (δ) .

A sketch of the equilibria is as follows. In the absence of simplicity biased consumers, firms share a mutual interest in confusing consumers and all firms adopt high complexity. This captures the equilibrium proposition in CZ's duopoly model of price-frame competition. There exists no incentive to undercut the obfuscation level of the rival and industry profit, to which each firm anticipates an equal share, is increasing in the proportion of uninformed consumers.

When simplicity biased consumers are introduced, competition is generated in obfuscation as firms seek to undercut the obfuscation of the rival. When the bias (δ) is small, this incentive to compete is dominated by the incentive to inject complexity to soften price competition. As the bias (δ) increases, the incentive to deviate from high complexity increases, leading firms to vary between high and low complexity. For a sufficiently strong bias (δ) , firms may be disciplined to low complexity, depending on model parameters.

3.4.1 Symmetric Equilibria

First consider pure strategy symmetric equilibria. Suppose that the combination $\{L, L\}$ is an equilibrium, producing a maximally transparent market.¹⁵ By definition there exists no profitable deviation. By symmetry it is necessary only to show that one firm does not have a profitable deviation such that $\pi_i\{L, L\} \geq \pi_i\{H, L\}$ where expected profit for firm i , given the complexity choice of each firm, takes the form $\pi_i\{k_i, k_j\}$. This condition requires:

$$\frac{1 - \mu_L}{1 - \mu} \geq \frac{1 - \delta^2 + \mu(1 + \delta)^2}{1 + \delta + \mu(1 - \delta)} \quad (3.7)$$

¹⁵Transparency does not require that all consumers are informed. Rather, it refers to the minimum level of market complexity that can be achieved.

Proposition 3. A low complexity equilibrium is more likely when:

- (i) The degree of simplicity bias is high.
- (ii) The natural level of consumer sophistication (μ_L) is low.

The intuition follows that an increase in the natural level of consumer sophistication engenders two effects. First, a higher level of market complexity is required to achieve the same level of consumer myopia. Second, for a fixed μ , the difference ($\mu_L - \mu$) is increasing in μ_L . This difference corresponds to the marginal reduction in consumer sophistication when one firm deviates to high complexity, which captures the effectiveness of obfuscation. When obfuscation is more effective in confusing consumers, the marginal benefit from obfuscation increases, which strengthens the incentive for a firm to obfuscate. Therefore, policies that stimulate the natural proportion of informed consumers may lead firms to increase their obfuscation.

The condition also makes clear that there exists parameter values for which a low complexity equilibrium can never be sustained, regardless of the degree of simplicity bias. In this case the additional demand from simplicity biased consumers is dominated by the loss arising from more fierce price competition that would prevail with a larger group of informed consumers. This insight offers an explanation for the persistence of highly complex offers in markets where consumers exhibit strong preferences for simplicity. The increase in the likelihood of a low complexity equilibrium when the degree of bias increases is also intuitive as a larger bias raises the penalty for a firm that deviates to high complexity.

Consider now the possibility of a symmetric pure strategy equilibrium in which both firms adopt high complexity, which requires $\pi_i\{H, H\} \geq \pi_i\{L, H\}$ and translates into the following condition:

$$\frac{\mu - \mu_H}{1 - \mu} \geq \delta \quad (3.8)$$

Proposition 4. A high complexity equilibrium is more likely when:

- (i) The intermediate level of consumer sophistication (μ) is high.
- (ii) The minimum level of consumer sophistication (μ_H) is low.
- (iii) The degree of simplicity bias is low.

Remark 2. A high complexity equilibrium is always sustained if $\mu - \mu_H > 1 - \mu$.¹⁶

¹⁶The remainder of the article focuses on the more interesting case where $\mu - \mu_H < 1 - \mu$. This is a reasonable condition, given that in practice there are likely to exist diminishing effects from obfuscation. That is, as more of the consumer population are confused by market complexity, confusing an additional person becomes increasingly difficult.

It is immediately clear that by neglecting consumers preferences for simple options ($\delta = 0$) we would conclude that firms should always maximally obfuscate. When the degree of simplicity bias is small, $\{H, H\}$ is sustained because the incentive to deviate is insufficient. As the bias increases the incentive for a firm to deviate strengthens, overturning the high complexity equilibrium providing $\mu - \mu_H < 1 - \mu$. The precise intuition is made clearer by rewriting (3.8) in the form:

$$(\mu - \mu_H) \geq \delta(1 - \mu) \quad (3.9)$$

The condition states that a high complexity equilibrium is sustained only when the marginal effect of obfuscation exceeds the value of simplicity biased consumers. An increase in μ also softens the condition for a high complexity equilibrium due to three unidirectional effects. First, when μ increases, a higher level of obfuscation is required to achieve the same level of consumer myopia. Second, when μ increases, the marginal effect of a second firm obfuscating ($\mu - \mu_H$) increases. Third, a larger proportion of informed consumers necessarily implies a smaller proportion of uninformed consumers of which a fraction are simplicity biased. These effects cumulatively stimulate the incentive to obfuscate when μ is larger.

It is also salient that an increase in μ increases the degree of simplicity bias required to overturn the high complexity equilibrium. Therefore, the influence of simplicity biases will be most pronounced in industries where consumer confusion is high. The intuition for a high complexity equilibrium becoming more likely when μ_H is small follows from the same arguments, approached from the opposite direction.

3.4.2 Asymmetric Equilibria

There exists parameter constellations for which the conditions for a symmetric pure strategy equilibrium do not hold. In this case there exists both an asymmetric pure strategy equilibrium if firms are able to successfully coordinate and a mixed strategy equilibrium in which firms randomise between high and low complexity.

For an asymmetric equilibrium we require two conditions to hold simultaneously: $\pi_i\{L, H\} \geq \pi_i\{H, H\}$ and $\pi_j\{L, H\} \geq \pi_j\{L, L\}$, reflecting the absence of a profitable deviation for both firms. The intuition follows that the presence of simplicity biased consumers may generate sufficient competition between firms in terms of obfuscation to overturn the high complexity equilibrium but fail to generate sufficient incentives to prevent firms adopting a strategy that incorporates elements of both high and low complexity.

3.4.3 Mixed Strategy Equilibrium

In practice, firms may be unable to coordinate, invoking a symmetric mixed strategy equilibrium. This requires that firms randomise between low and high complexity according to a probability profile. Let $\lambda(L)$ denote the probability with which firm j selects low complexity and correspondingly let $\lambda(H) = 1 - \lambda(L)$ be the probability that firm j selects high complexity. Similar to the pricing stage, the optimal probability profile must make the firm indifferent over the pure strategies, which generates the following indifference condition:

$$\lambda(L) \cdot \pi_i\{L, L\} + \lambda(H) \cdot \pi_i\{L, H\} = \lambda(L) \cdot \pi_i\{H, L\} + \lambda(H) \cdot \pi_i\{H, H\}$$

Insert the relevant payoffs and rearrange to yield the optimal probability with which low complexity is chosen:

$$\lambda(L) = \frac{\delta - \delta\mu + \mu_H - \mu}{\delta - \delta\mu + \mu_H - \mu - 1 + \mu_L + \left[\frac{(1-\mu)(1-\delta^2 + \mu(1+2\delta+\delta^2))}{1+\delta+\mu(1-\delta)} \right]} \quad (3.10)$$

To conform with the properties of probability we require $\lambda(L) \in [0, 1]$. For $\lambda(L) \leq 0$, the firm places zero probability on low complexity (High complexity equilibrium). For $\lambda(L) \geq 1$, the firm places zero probability on high complexity (Low complexity equilibrium). Consistent with the comparative statics of the pure strategy equilibria, the probability that low complexity is chosen increases with the degree of simplicity bias. The complete proof is provided in the appendix.

Result 3. *In mixed strategy equilibrium the probability placed on low complexity by each firm is increasing in the degree of simplicity bias.*

The primary interpretation of this mixed strategy is that firms change the complexity of their products or promotions over time. This differs from existing explanations for firms to change the presentation of their product that are based on preventing consumers from learning prices because in these models firms change only to a different promotion format of the same complexity level (Carlin & Manso, 2011). In contrast, our model predicts that firms will also vary the complexity level of their product.

Alternatively the mixed strategy can be interpreted as firms having access to a range of product dimensions that influence their complexity. This captures the essence of Spiegelger's (2006) contingency-based pricing model except firms now vary the complexity of each dimension such as a simple website complicated by excessive terms and conditions if consumers only consider one of the dimensions. This line of thought offers an explanation for firms to offer products with both simple and complex dimensions.

The following example illustrates the results. Suppose the natural proportion of informed consumers is $\mu_L = 0.75$ but obfuscation reduces consumer sophistication at a constant rate such that $\mu = 0.5$ and $\mu_H = 0.25$. The resulting probability distribution between low and high complexity $\lambda(L)$, for varying degrees of simplicity bias, is given in Figure 3.5.

When fewer than half of the uninformed consumers are simplicity biased, the incentive to compete in obfuscation is insufficient to break the high complexity equilibrium and $\lambda(L) = 0$. As the simplicity bias strengthens, the incentive to undercut the complexity of the rival increases and firms randomise between high and low complexity according to the probability measure $\lambda(L)$, which is increasing in the degree of bias. In the limiting case where $\delta = 1$, firms randomise equally between high and low complexity. Interestingly for this combination of parameters there exists no equilibrium in which firms are disciplined to low complexity.




Figure 3.5: Equilibrium Obfuscation I




Figure 3.6: Equilibrium Obfuscation II

The presence of simplicity biased consumers is therefore more likely to generate an intermediate level of market transparency than a fully transparent market. If the parameters of the model are reduced to $\mu_L = 0.45$, $\mu = 0.3$, $\mu_H = 0.15$, the natural level of consumer sophistication is lower and the marginal effect of obfuscation is weaker. The high complexity equilibrium is now overturned by $\delta > 0.21$ and a fully transparent market is achieved whenever $\delta > 0.66$, as illustrated in Figure 3.6.

3.5 Welfare Analysis

The presence of simplicity biased consumers generates two effects on industry profit. Firstly, firms seek to capitalise on consumers' preferences by bundling simple products with high prices, which leads simplicity biased consumers to pay the highest (average) price and increases firm profit. We term this the Direct Effect of the simplicity bias. Secondly, simplicity biased consumers engender a positive externality by incentivising firms to reduce their obfuscation. This leads to lower market complexity and a corresponding increase in the proportion of informed consumers. As a consequence firms compete more fiercely in prices and firm profit declines. We term this the Externality Effect. Therefore, the incidence of simplicity biased demand on industry profit and consumer welfare depends critically on which effect dominates.

Expected profit for each firm can be written in the form:

$$E[\pi_i] = 0.5 \left[1 - \mu_H + \lambda(L) \cdot \left(\frac{(1 - \mu)(1 - \delta^2 + \mu + 2\delta\mu + \delta^2\mu)}{1 + \delta + \mu - \delta\mu} - 1 + \mu_H \right) \right] \quad (3.11)$$

Proposition 5. The effect of the simplicity bias on firm profit is:

1. Profit is always reduced if the simplicity bias leads to a low complexity equilibrium.
2. If a low complexity equilibrium is not obtained:
 - (i) For $\mu + 2\mu_H \leq 1$, an increase in the bias always reduces profit.
 - (ii) For $\mu + 2\mu_H > 1$, an increase in the bias may increase or decrease profit depending on parameter values for consumer sophistication.

In a low complexity equilibrium both firms are equally complex and the proportion of informed buyers is maximised, which leads to maximum price competition.

If a mixed strategy equilibrium emerges in obfuscation, firm profit declines in the degree of bias whenever $\mu + 2\mu_H \leq 1$ because the externality effect dominates the direct effect. This level of consumer sophistication fits well with many retail markets. For example, the Advertising Standards Authority report that only 23% of consumers could identify the total cost of broadband advertisements on first viewing.¹⁷

In markets where consumers are more informed, however, such that $\mu + 2\mu_H > 1$, we are confined to numerical analysis because the absolute and relative values of the parameters are

¹⁷www.asa.org.uk/News-resources/Media-Centre/2016/We-confirm-tougher-approach-to-broadband-price-claims-in-ads. Accessed 04/05/16.

interactive in determining the magnitude of the direct and externality effects. Table 3.1 outlines 5 numerical examples for the effect of an increase in the degree of bias on profit.¹⁸

μ_L	μ	μ_H	$\frac{\partial \pi_i}{\partial \delta}$
0.75	0.5	0.25	< 0
0.9	0.6	0.4	< 0
0.8	0.7	0.6	-
0.5	0.4	0.3	< 0
0.6	0.4	0.2	< 0

Table 3.1: Numerical Profit Examples

These results are especially interesting for two reasons. First, they imply that consumers who select the simplest option will pay the highest price but firms may earn lower profit as the degree of bias increases. This paradox arises from the reduction in obfuscation that biased consumers generate, which increases the proportion of informed buyers and stimulates price competition. Second, in our framework there exists parameter values where consumer welfare is lower when obfuscation is lower and more consumers are informed. Therefore, any case for regulatory intervention should not be motivated solely by high market complexity.

3.6 Discussion

This article develops a model of competition that explicitly incorporates consumers' preferences for simple products. At a time when regulatory intervention is increasingly motivated by imperfect consumer behaviour, we document a mechanism by which non-standard decision-making generates novel dimensions of competition that stimulate market transparency and may improve consumer welfare. Furthermore, regulation in this setting may reduce market transparency and fail to benefit consumers.

Consistent with mainstream literature, the model predicts that firms use obfuscation to escape competition: Obfuscation leads to fewer informed consumers, which softens price competition and insulates industry profit. This crystallises the collective incentive for firms to inject spurious complexity.

At the firm level, however, the correlation between obfuscation and prices depends on the degree of dispersion in firms' complexity. When firms select the same complexity level, sellers select a price randomly from identical non-degenerate distributions. There exists no correlation between

¹⁸ – denotes a non-monotonic relationship created by the counter-veiling direct and externality effects. In the appendix we illustrate that for a wide range of parameter values, profit is decreasing in the degree of bias.

a firm's complexity level and price. In contrast, when firms adopt different complexity levels, a negative correlation is generated by the simplicity bias. The firm with the lowest complexity level secures a larger segment of captive consumers, which generates a weaker incentive to engage in price competition. The model also shows that more complex firms persistently offer fluctuating 'sale' prices whilst less complex firms vary between 'sale' prices and the RRP.

The main parameter of interest in the model is δ , which captures the degree of simplicity bias amongst consumers who cannot comprehend prices. The model shows that the simplicity bias is effective in overturning the high complexity equilibrium, providing the degree of bias is sufficiently strong and $2\mu - \mu_H < 1$. The ability of simplicity biased consumers to discipline firms to low complexity, however, is dependent on model parameters. A transparent market is most likely when the marginal effect of complexity on consumer sophistication is small, the natural level of informed consumers is low and the degree of simplicity bias is high.

The incidence of the simplicity bias on consumer welfare is multifaceted. When consumers that are impaired by complexity buy from the least complex seller, they pay the highest price. However, they exert a positive externality on the market that suppresses obfuscation. Transparency increases and more consumers are able to identify the lowest priced seller. Price competition subsequently intensifies and consumer welfare may increase at the aggregate level. Therefore, consumer protection policies designed to debias consumers from selecting simple products will reduce market transparency and may reduce consumer welfare.

The observation that firms may choose not to obfuscate is also consistent with recent empirical evidence. Miravete (2013) finds no evidence of firms responding to competition by increasing the foginess of their pricing structures in the Telecommunications industry. Therefore, the idea that regulation is required to limit obfuscation may be misplaced and we offer a direct explanation for firms not obfuscating.

The effect of regulatory intervention to improve consumer sophistication depends on the existing state of the market and the precise incidence of policy on model parameters. If firms are initially operating at high complexity, a regulator could provide decision-making aids to increase the proportion of informed consumers. Price competition would intensify and consumer welfare would increase. Furthermore, when the effectiveness of obfuscation is smaller, a firm is more likely to deviate to low complexity because the relative merits of obfuscation decline. Therefore, regulation may induce a positive multiplied effect. However, if firms are initially operating at low complexity, measures to improve consumer sophistication may lead firms to increase their obfuscation, which is consistent with Gu and Wenzel (2014).

Policy is therefore most effective in industries where the existing level of market complexity is high and the degree of simplicity bias exhibited by uninformed consumers is strong. These results also provoke an empirical enquiry to investigate any systematic characteristics of informed,

simplicity biased and uninformed consumers to assess the distributional impact of the simplicity bias when authorities overweight the welfare of particular groups.

In some environments consumers' perception of the complexity associated with a firm's product may be fixed and independent of obfuscation. For example, online retailers that deliver to a customer may be perceived as intrinsically more simple than collection-only equivalents. *Ceteris paribus* delivered products should exhibit greater price dispersion and be frequently priced at the RRP. In-store prices, which are associated with a larger degree of complexity due to locating the store and the product, should always be priced below the RRP and exhibit lower price dispersion.¹⁹ This result is consistent with Narasimhan (1988) and Gu and Wenzel (2014) but with an interpretation of heterogeneity in the convenience of a suppliers service rather than market prominence. Furthermore, these results are empirically testable with relatively low data intensity.

The model raises questions regarding why consumers continue to select products on the basis of simplicity if they are more expensive. This can be explained in at least three ways. First, in many retail environments it is reasonable to assume that consumers do not learn from their mistakes or are not aware that a suboptimal choice has been made. For example, consumers are often unaware of their energy tariff or pension investment scheme (Wilson and Waddams-Price, 2010).²⁰ Second, the model assumes that consumers do not derive utility from consuming an option that is simple. In practice, consumers may derive non-monetary gains from the purchase of a simple option. This also links to concerns for price fairness, perceived firm transparency and the overestimation of more complex pricing structures (Homburg et al., 2014). Third, consumers could act in concert by purchasing from the least complex seller to prevent firms from obfuscating if a low complexity equilibrium is feasible. It is likely, however, that a rational consumer might defect to purchase from the high complexity firm to secure a lower price, if a low complexity equilibrium cannot be obtained. This final point, however, offers a role for policy if market authorities are able to encourage collective decision-making by consumers or achieve this effect artificially using government purchases.

The model could also be written without a behavioural foundation by introducing asymmetric search costs across the population of consumers. However, in many markets search costs are small but consumers continue to exhibit a preference for simplicity. This validates the 'all-or-nothing' search approach, which is widely established in the literature. Further research to understand the exact origins of the simplicity bias and the variation in the strength of consumers' simplicity biases across industries, however, may be particularly interesting. Specifically, the model predicts that

¹⁹Heterogeneity in price dispersion arises because both firms draw prices from distributions with the same lower bound but with the more complex firm pricing strictly below the reservation price, whilst the distribution of the less complex rival includes the reservation price.

²⁰In the model the supports of the price distributions are also nearly identical when firms select different complexity levels, which makes this explanation particularly plausible.

in markets where consumers exhibit a stronger preference for less complex products, obfuscation will be lower and prices may be lower, which is empirically testable.

The model also suggests a novel avenue for false advertising. In the model, firms balance the incentive to inject spurious complexity with the incentive to undercut the complexity of the rival and we assumed that complicated offers, which reduce consumer sophistication, are correctly perceived as being the most complex option by consumers. However, it may be possible for a firm to engineer a marketing strategy that stimulates market complexity whilst appearing simple to consumers. This satisfies the competing incentives simultaneously. Critically, this does not require that firms are able to manipulate consumers' preferences.²¹

This investigation also highlights an interesting possible explanation for multi-brand retailing in which firms owned by a single parent organisation may follow different promotion strategies. Ireland (2007) demonstrates that when a firm sells under multiple brand names and consumer search is costly, the optimal pricing strategy requires prices to be identical across outlets. The intuition follows that consumers purchase from the lowest priced supplier in their sample. If a consumer samples multiple stores that are all owned by a single firm, any dispersion in prices constitutes an unnecessary loss to the firm. In this sense, multi-brand retailing exists only to create an environment of illusory competition and exploit consumers' search costs. However, in practice, brands owned by a single organisation often differ in their promotion styles.

The direction from our model suggests that a firm might use one brand to stimulate consumer confusion and another to attract simplicity biased consumers. The high complexity outlet shifts consumers towards the simple outlet, which is subsequently able to charge a higher price. The reason obfuscation may vary across outlets whilst prices may not emanates from the fact that pricing only influences firm demand with no effect on consumer sophistication. In contrast, obfuscation influences both firm demand and the composition of the consumer population in terms of the proportion of informed and uninformed consumers.

This intuition may also explain why a firm chooses to operate under multiple brand names, which impedes consumer comparisons, but also own a price comparison website that simplifies decision-making. For example, Confused.com, a leading UK price comparison tool, also operates several insurance brands featured on the comparison website.²² Future research should examine these incentives to determine how such equilibria might look.

²¹The connection with false advertising is formalised in Chapter 4.

²²For further information see www.confused.com/about-us/how-we-operate.

3.6.1 The Role of Timing

In the model, firms select an obfuscation level before setting prices because in practice there exists significant lag periods involved in augmenting the complexity of a firm's product, whereas prices are more instantaneously adjustable. In Chapter 4 we also show that many of the results are preserved when firms select price-obfuscation pairs simultaneously, amongst other refinements.

In the restatement, the presence of simplicity biased consumers always improves consumer welfare and reduces the propensity of firms to obfuscate. However, in contrast with the sequential game, consumers who select the least complex option benefit directly from their simplicity bias because the least complex firm charges the lowest price. The intuition is that firms are unable to condition their price on the obfuscation of the rival. This leads a firm with a high price to obfuscate more than a firm with a low price, to conceal their lack of competitiveness.

3.6.2 Divergent Consumer Preferences

When preferences are identical, all consumers agree which firm offers the simplest product. However, in practice consumers may differ in their perceptions. To address this it is sufficient to recognise that divergences in consumers' preferences do not alter the model significantly, providing more consumers prefer the 'simple' option. The reason is that when a minority prefer the 'complex option', some of the consumers selecting the 'simple option' are neutralised with the only effect being a reduction in δ .

Consider the following example. There are 5 uninformed consumers; 2 purchase randomly but 3 are simplicity biased ($\delta = 0.6$). The main model requires that these 3 consumers purchase from the firm with the least complex offer. However, providing at least 2 select the simple offer, the results hold with a reduced δ . If 2 consumers prefer the simple option whilst 1 prefers the complex option, the model retains the core framework with $\delta = 0.2$. That is, 4 consumers select randomly and 1 is simplicity biased.

3.6.3 Costly Obfuscation

The primary analysis assumes obfuscation is costless for firms, which is restrictive. In practice, engineering a complex promotion format requires effort. Furthermore, secondary costs are likely to arise as consumers make greater use of customer services to clarify their purchase. To capture this, suppose obfuscation is costly in the sense that firms incur a positive fixed cost, $c > 0$, if a firm chooses high complexity.

The conditions for low and high complexity equilibria respectively become:

$$\frac{1 - \mu_L}{1 - \mu} \geq \frac{1 - \delta^2 + \mu(1 + 2\delta + \delta^2)}{1 + \delta + \mu(1 - \delta)} - \frac{2c}{r(1 - \mu)} \quad (3.12)$$

$$\frac{\mu - \mu_H}{1 - \mu} - \frac{2c}{r(1 - \mu)} \geq \delta \quad (3.13)$$

Result 4. *If obfuscation is costly for firms:*

1. *Increasing the direct cost of obfuscation leads to lower obfuscation by firms.*
2. *For $c \geq \frac{1}{2} \cdot \left[\frac{(1-\mu)(1-\delta^2+\mu(1+2\delta+\delta^2))}{1+\delta+\mu-\delta\mu} - (1 - \mu_L) \right]$, all firms are disciplined to low complexity in stage 1.*

Result 4 is salient for competition authorities because increasing the cost of obfuscation directly reduces firms' incentive to inject complexity and improves market transparency. Therefore, policies to suppress obfuscation should stimulate this direct cost when traditional measures such as decision-making aids and campaigns to encourage consumer engagement are impotent, as firms respond with higher levels of obfuscation.

This line of thought creates a role for ex-post regulation in the spirit of Acutt et al. (2001). They show that firms may be incentivised to operate competitively by a regulator that establishes a credible reputation for punishing malpractice whenever it occurs. In the present setting, regulators could introduce penalties for firms that are identified to have obfuscated, which targets firms' incentives directly.

3.6.4 Further Extensions

The generality of the model offers numerous avenues for extensions. In particular, the main model assumes that each firm is equally effective in confusing consumers. In practice, one firm may be more effective at stimulating market complexity. Furthermore, we assumed that firms are equally prominent in the market, leading uninformed consumers to exhibit no preference between the two firms when identical complexity levels are chosen. In reality, market prominence plays a significant role in consumer decision-making.²³

²³See De los Santos et al. (2012) for a recent empirical study of consumer purchasing behaviour in which 69% of consumers immediately purchase from the market leader in the online market for books.

Gu and Wenzel (2014) examine obfuscation behaviour when prominence varies across firms due to exogenous factors. Incorporating their exogenous prominence, which captures brand loyalty amongst consumers, with the endogenous prominence in our model, which captures consumers' preferences for simplicity, may yield interesting results. Future studies might also endogenise the value of δ , perhaps as a function of firms' advertising efforts.

The model does not consider participation effects whereby a potential customer is perturbed by market complexity and does not purchase from any firm because in many retail environments the consumer is compelled to purchase the underlying product. This might include domestic energy, mobile phone contracts and a bank. The introduction of complexity-induced non-participation effects would soften the incentive for a firm to obfuscate further.

It would also be beneficial to examine an oligopoly version of the model we analyse in this article to determine how firms respond when market concentration decreases. Carlin (2009), amongst others, demonstrates that firms respond to market entry by increasing their obfuscation. In the present model it is not clear whether this will be the case as competition between firms in obfuscation might intensify. To achieve this, however, requires a more detailed understanding of consumers' choice procedures.

3.7 Conclusion

In this article we show that the incentives for a firm to obfuscate are significantly weaker when consumers are biased towards simple products. In particular, consumers who buy from the least complex firm pay the highest price but their presence generates a positive externality by reducing the obfuscation of firms. As a result, more consumers are able to identify the lowest priced firm, which stimulates price competition. Therefore, policies that discourage consumers from selecting simple products lead to more complex market interactions for buyers and can lead to higher prices. In addition, policies to improve consumer sophistication may also lead firms to increase their obfuscation.

The model also provides an explanation for firms to vary both their offer and the complexity of their offer over time. This fits well with many retail markets. Furthermore, firms that offer simple products vary between sale and non-sale prices but more complex firms continually randomise over sale prices, with a lower average price. The article also highlights the incentives for a firm to engineer a promotion format that simultaneously limits comparability for consumers but appears simple. Finally, this article shows that the introduction of simplicity biased consumers offers new avenues of study regarding multi-brand retailing and information intermediaries.

3.8 Appendix A: Main Proofs

Proof of Proposition 1. Selecting a price deterministically makes the firm vulnerable to the usual undercutting principle as in Varian (1980). Suppose that there exists a price $0 \leq p' \leq r$ for firm i that is a pure strategy Nash equilibrium. By definition p' must be a best response to any strategy of the rival. The rival could select a price $p' - \epsilon$, $\epsilon > 0$, which discontinuously reduces the payoff of firm i , contradicting p' as a pure strategy equilibrium. This holds across the domain of optimal equilibrium prices. ■

Proof of Proposition 2. The upper bound of the price distribution must be the consumers' reservation price. Any price above this level yields zero demand. Any maximum price, $\bar{p} < r$, only attracts uninformed consumers who would purchase at r . Raising the price to r yields a strict improvement in firm profit, contradicting \bar{p} being the upper bound of the support.

The presence of uninformed consumers guarantees that every firm makes positive profit. At $p_i \leq 0$ the firm is guaranteed a segment of uninformed consumers and providing $p_i < p_j$ the entire population of informed consumers but a non-positive profit is received. Increasing the price to $p_i \in (0, r]$ continues to guarantee the custom of uninformed consumers and a possibility of informed consumers, which guarantees positive profit. ■

Proof of Result 1.

Claim. In the symmetric pricing case firms draw prices randomly from a continuous, atomless and well defined probability distribution function with support $\left[\frac{1-\mu_L}{1+\mu_L}r, r\right]$. The distribution function is also increasing in the firms price, p .

Proof. The indifference condition for the high complexity equilibrium is given by:

$$\frac{(1-\mu_L)}{2} \cdot r = \left[\frac{1-\mu_L}{2} + \mu_L [1 - F^L(p)] \right] p$$

Solving for $F^L(p)$ yields the distribution:

$$F^L(p) = 1 - \frac{(1-\mu_L)}{2\mu_L} \left[\frac{r}{p} - 1 \right]$$

The lower bound of the pricing distribution is given at $F^L(p) = 0$, which attains at $\frac{1-\mu_L}{1+\mu_L}r$. Inserting the minimum price into the distribution function above confirms the solution. Furthermore, $F^L(r) = 1$, confirming that the firm places zero probability on prices exceeding the consumers reservation threshold. Differentiating the distribution with respect to price, p , yields:

$$\frac{\partial F^L(p)}{\partial p} = + \frac{(1 - \mu_L) r}{2\mu_L p^2} > 0$$

The function must also be continuous by the usual arguments. This confirms the legitimacy of the candidate solution. In the symmetric pricing case where firms adopt high complexity, the proof of the results follows identically. ■

Proof of Result 2. The proof is separated into two claims.

Claim. In the asymmetric pricing case firms draw prices randomly from different continuous probability distribution functions with identical lower bounds. The only atom exists in the distribution of the least complex firm at the reservation price, which the more complex rival always prices below. Both price distributions are increasing in p .

Proof. Suppose throughout that firm i is less complex than firm j . The indifference condition for the least complex firm (i) is written:

$$\frac{(1 + \delta)(1 - \mu)r}{2} = \left[\frac{(1 + \delta)(1 - \mu)}{2} + \mu[1 - F_j(p)] \right] p$$

Set $F_j(p) = 0$ to yield the minimum price (\underline{p}):

$$\frac{(1 + \delta)(1 - \mu)r}{1 + \delta + \mu - \delta\mu} = \underline{p}$$

Solving the indifference condition for the distribution function yields:

$$F_j(p) = 1 - \frac{(1 + \delta)(1 - \mu)}{2\mu} \left[\frac{r}{p} - 1 \right]$$

Insert the minimum price into the distribution function to demonstrate $F_j(p) = 0$, confirming zero probability is placed on prices below the lower bound. Differentiate the distribution function to yield:

$$\frac{\partial F_j(p)}{\partial p} = + \frac{(1 + \delta)(1 - \mu)}{2} \cdot \frac{r}{p^2} \geq 0$$

The indifference condition for the more complex firm (j), taken at the lower bound of the price distribution, is written:

$$\frac{(1 + \delta)(1 - \mu)r}{1 + \delta + \mu(1 - \delta)} \cdot \left(\frac{(1 - \delta)(1 - \mu)}{2} + \mu \right) = \left(\frac{(1 - \delta)(1 - \mu)}{2} + \mu[1 - F_i(p)] \right) p$$

Solve for the price distribution:

$$F_i(p) = 1 - \frac{(1-\mu)}{2\mu} \left[\frac{(1+\delta)(1-\delta+\mu+\delta\mu)r}{(1+\delta+\mu(1-\delta))p} - (1-\delta) \right]$$

Verify the value of the mass point by solving $F_i(r)$:

$$F_i(r) = 1 - \frac{2\delta(1-\mu)}{1+\delta+\mu(1-\delta)}$$

Verify the lower bound of the price distribution where $\mathbf{p} = \frac{(1+\delta)(1-\mu)r}{1+\delta+\mu(1-\delta)}$. (These steps need not be published.)

$$\begin{aligned} F_i(\mathbf{p}) &= 1 - \frac{(1-\mu)}{2\mu} \left[\frac{(1+\delta)(1-\delta+\mu+\delta\mu)r}{(1+\delta+\mu(1-\delta))\mathbf{p}} - (1-\delta) \right] \\ F_i(\mathbf{p}) &= 1 - \frac{(1-\mu)}{2\mu} \left[\frac{(1+\delta)(1-\delta+\mu+\delta\mu)r(1+\delta+\mu-\delta\mu)}{(1+\delta+\mu(1-\delta))(1+\delta)(1-\mu)r} - (1-\delta) \right] \\ F_i(\mathbf{p}) &= 1 - \frac{(1-\mu)}{2\mu} \left[\frac{(1-\delta+\mu+\delta\mu)}{(1-\mu)} - (1-\delta) \right] \\ F_i(\mathbf{p}) &= 1 - \frac{1}{2\mu} [(1-\delta+\mu+\delta\mu) - (1-\delta-\mu+\delta\mu)] = 1 - \frac{2\mu}{2\mu} = 0 \end{aligned}$$

Differentiate the distribution function with respect to p to confirm $\frac{\partial F_i(p)}{\partial p} \geq 0$.

$$\frac{\partial F_i(p)}{\partial p} = + \frac{(1-\mu)}{2} \cdot \frac{(1+\delta)(1-\delta+\mu+\delta\mu)r}{(1+\delta+\mu(1-\delta))p^2} \geq 0$$

Claim. The least complex firm secures a larger payoff than the rival.

Proof. When firms differ in their obfuscation level, the payoff for the least complex firm is written $\frac{(1+\delta)(1-\mu)}{2}$ and the payoff for the most complex firm can be written as $p_j \rightarrow r$ following Narasimhan (1988): $\left[\frac{(1-\delta)(1-\mu)r}{2} + \frac{2\delta\mu(1-\mu)r}{1+\delta+\mu-\delta\mu} \right]$. It is now possible to show that for any $\mu \in (0, 1)$ the payoff to the least complex firm is larger:

$$\frac{(1+\delta)(1-\mu)}{2} > \frac{(1-\delta)(1-\mu)}{2} + \frac{2\delta\mu(1-\mu)}{1+\delta+\mu-\delta\mu}$$

This condition simplifies to: $(1 - \mu)(1 + \delta) > 0$, which is satisfied for any $\delta \in (0, 1]$ and any $\mu \in (0, 1)$. Therefore, the least complex firm receives a larger payoff than the most complex firm, which underpins the incentive for sellers to compete in obfuscation. ■

Proof of Proposition 3. Define $\phi := \frac{1 - \mu_L}{1 - \mu}$. The proportion of informed consumers is decreasing in obfuscation such that $\mu_L > \mu$. This implies $\phi \in (0, 1)$. Rewriting the condition for a low complexity equilibrium in terms of μ yields:

$$\frac{-1 + \phi + \delta^2 + \delta\phi}{1 - \phi + 2\delta + \delta\phi + \delta^2} \geq \mu \quad (3.14)$$

Differentiate the left hand side (L.H.S.) of (3.14) with respect to δ :

$$\frac{\partial}{\partial \delta} \left(\frac{-1 + \phi + \delta^2 + \delta\phi}{1 - \phi + 2\delta + \delta\phi + \delta^2} \right) = \frac{4\delta - 4\delta\phi + 2 - 2\phi^2 + 2\delta^2}{(1 - \phi + 2\delta + \delta\phi + \delta^2)^2} \geq 0$$

The L.H.S. is increasing in δ , which makes the condition less stringent. The second element is visible by rewriting (3.7) in terms of μ_L :

$$1 - \frac{(1 - \mu)(1 - \delta^2 + \mu(1 + 2\delta + \delta^2))}{1 + \delta + \mu - \delta\mu} \geq \mu_L$$

An increase in μ_L makes the condition more stringent. ■

Proof of Proposition 4. The condition for a high complexity equilibrium is given by:

$$\frac{\mu - \mu_H}{1 - \mu} \geq \delta \quad (3.15)$$

Differentiate the L.H.S. of (3.15) with respect to μ and μ_H to yield:

$$\frac{\partial}{\partial \mu} \left(\frac{\mu - \mu_H}{1 - \mu} \right) = \frac{1 - \mu_H}{(1 - \mu)^2} \geq 0 \quad \frac{\partial}{\partial \mu_H} \left(\frac{\mu - \mu_H}{1 - \mu} \right) = -\frac{1}{1 - \mu} \leq 0$$

■

Proof of Result 3. The proof that $\frac{\partial \lambda(L)}{\partial \delta} > 0$ in mixed strategy equilibrium proceeds in steps. *Step 1:* Write $\lambda(L)$ in the form:

$$\frac{\delta - \delta\mu + \mu_H - \mu}{\delta - \delta\mu + \mu_H - \mu + \mu_L - 1 + \frac{(1 - \mu)(1 - \delta^2 + \mu - 2\delta\mu + \delta^2\mu)}{1 + \delta + \mu - \delta\mu}}$$

In mixed strategy equilibrium: $\delta(1-\mu) > \mu - \mu_H$ and $\frac{(1-\mu_L)}{(1-\mu)} < \frac{1-\delta^2+\mu+2\delta\mu+\delta^2\mu}{1+\delta+\mu-\delta\mu}$, which implies both the numerator and denominator are positive. $\lambda(L)$ is a probability, implying $\lambda(L) \leq 1$ in mixed strategy equilibrium. We now ask the question, does increasing δ make the condition $\lambda(L) \leq 1$ more or less likely? If the solution is less likely, $\lambda(L)$ must be increasing in δ . Therefore, rewrite the condition:

$$\frac{\delta - \delta\mu + \mu_H - \mu}{\delta - \delta\mu + \mu_H - \mu + \mu_L - 1 + \frac{(1-\mu)(1-\delta^2+\mu+2\delta\mu+\delta^2\mu)}{1+\delta+\mu-\delta\mu}} \leq 1$$

in the form:

$$\frac{1 - \mu_L}{1 - \mu} \leq \frac{(1 - \delta^2 + \mu + 2\delta\mu + \delta^2\mu)}{1 + \delta + \mu - \delta\mu}$$

Step 2: By construction $\mu_L > \mu$. Therefore, $\frac{1-\mu_L}{1-\mu} < 1$. The proof now follows the structure of the proof of Proposition 3. Define $\phi := \frac{1-\mu_L}{1-\mu}$, where ϕ is independent of δ , and rewrite the above condition:

$$\phi \leq \frac{(1 - \delta^2 + \mu - 2\delta\mu + \delta^2\mu)}{1 + \delta + \mu - \delta\mu}$$

Step 3: Rearrange and simplify in terms of μ :

$$\mu \geq \frac{\delta^2 + \phi + \delta\phi - 1}{1 + 2\delta + \delta^2 - \phi + \delta\phi}$$

This is an equivalent expression to $\lambda(L) \leq 1$. It now follows from the proof of Proposition 3 that the R.H.S. is increasing in δ for all $\phi \in (0, 1)$, which implies that the condition $\lambda(L) \leq 1$ becomes more stringent as δ increases. To see this differentiate the R.H.S. of the above expression with respect to δ to yield:

$$\frac{4\delta - 4\delta\phi + 2 - 2\phi^2 + 2\delta^2}{(1 + 2\delta + \delta^2 - \phi + \delta\phi)^2} \geq 0$$

The derivative is positive, confirming that $\frac{\partial\lambda(L)}{\partial\delta} > 0$ in the relevant range. The reason for solving in this way is that in some range $\pi_i\{H, L\}$ is increasing in δ because the mass point of the rival increases, leading the more complex firm to yield a higher payoff. However, in this range $\lambda(L) = 0$. ■

Proof of Proposition 5. We separate the proof into 3 parts. First recognise that in the pure strategy equilibrium for obfuscation, the payoff when firms adopt low complexity is lower than the payoff when firms adopt high complexity because $\mu_L > \mu_H$. Therefore, $(1 - \mu_L) < (1 - \mu_H)$ and profit is lower in a low complexity equilibrium than a high complexity equilibrium.

Second, we show that for $\mu + 2\mu_H \leq 1$, profit is decreasing in δ . For this proof let $\lambda(L) = \lambda$ and $\lambda(H) = 1 - \lambda$.

From Result 3 we know that $\lambda(L)$ is increasing in δ . The expected payoff for a firm is given by:

$$\lambda^2 \cdot \pi_i\{L, L\} + \lambda \cdot (1 - \lambda) \cdot \pi_i\{L, H\} + (1 - \lambda)\lambda \cdot \pi_i\{H, L\} + (1 - \lambda)^2 \cdot \pi_i\{H, H\} \quad (3.16)$$

We know by definition of a mixed strategy equilibrium that for all $\lambda \in [0, 1]$:

$$\lambda \cdot \pi_i\{L, L\} + (1 - \lambda) \cdot \pi_i\{L, H\} = \lambda \cdot \pi_i\{H, L\} + (1 - \lambda) \cdot \pi_i\{H, H\}$$

Multiply by λ to achieve:

$$\lambda^2 \cdot \pi_i\{L, L\} + \lambda(1 - \lambda) \cdot \pi_i\{L, H\} = \lambda^2 \cdot \pi_i\{H, L\} + \lambda(1 - \lambda) \cdot \pi_i\{H, H\}$$

Substitute into (3.16) and rearrange to yield:

$$\pi_i\{H, H\} + \lambda \cdot (\pi_i\{H, L\} - \pi_i\{H, H\}) \quad (3.17)$$

Differentiate w.r.t δ :

$$\frac{\partial \lambda}{\partial \delta} \cdot [\pi_i\{H, L\} - \pi_i\{H, H\}] + \lambda \cdot \frac{\partial \pi_i\{H, L\}}{\partial \delta} \quad (3.18)$$

$\mu + 2\mu_H \leq 1$ guarantees that $\pi_i\{H, L\}$ is decreasing in δ . Therefore, the R.H.S is negative. The L.H.S. is also negative because at the critical point $\delta = \frac{\mu - \mu_H}{1 - \mu}$ for which $\lambda(L) = 0$ exactly, $\pi_i\{H, H\} = \pi_i\{L, H\}$. We also know from Remark 1 that $\pi_i\{L, H\} > \pi_i\{H, L\}$. Therefore, $\pi_i\{H, L\} < \pi_i\{H, H\}$ at the point $\lambda(L) = 0$. Furthermore, $\pi_i\{H, H\}$ is independent of δ but $\pi_i\{H, L\}$ is decreasing in δ for $\mu + 2\mu_H \leq 1$ such that $[\pi_i\{H, L\} - \pi_i\{H, H\}]$ is negative in MSNE.

Third, for $\mu + 2\mu_H > 1$ we cannot guarantee $\pi_i\{H, L\}$ is decreasing in δ in MSNE, which forces the numerical analysis provided in the main text.

The proof that $\frac{\pi_i\{H,L\}}{\partial\delta} < 0$ for $\mu + 2\mu_H \leq 1$ proceeds in steps. (Not for Publication).

Step 1: Write $\pi_i\{H, L\}$ in the form: $\frac{(1-\mu)(1-\delta^2+\mu+2\delta\mu+\delta^2\mu)}{2(1+\delta+\mu-\delta\mu)}$

Step 2: Rewrite $\pi_i\{H, L\}$ as: $\frac{(1-\mu)}{2} \left[\frac{(1-\delta^2+\mu+2\delta\mu+\delta^2\mu)}{(1+\delta+\mu-\delta\mu)} - 1 + 1 \right]$

Step 3: Simplify to obtain: $\frac{(1-\mu)}{2} \left[\frac{(-\delta^2+\delta^2\mu+3\delta\mu-\delta)}{(1+\delta+\mu-\delta\mu)} + 1 \right]$

Step 4: Differentiate w.r.t δ to obtain: $\frac{(1-\mu)}{2} \cdot \left(\frac{(1-\mu)(-\delta^2+\delta^2\mu)+(1+\mu)(-2\delta+2\delta\mu+3\mu-1)}{(1+\delta+\mu-\delta\mu)^2} \right)$

$(-\delta^2 + \delta^2\mu) < 0$ because $\mu \in (0, 1)$. The derivative is always negative if: $2\delta(1 - \mu) > 3\mu - 1$. However, in MSNE $\delta > \frac{\mu - \mu_H}{1 - \mu}$, otherwise $\lambda(L) = 0$. Therefore we can relax this condition to: $2 \cdot \frac{\mu - \mu_H}{1 - \mu} (1 - \mu) > 3\mu - 1 \implies 1 > 2\mu_H + \mu$. ■

Proof of Result 4. The condition for a low complexity equilibrium with a fixed cost c whenever a firm selects high complexity, is written:

$$\frac{1 - \mu_L}{2} r \geq \frac{(1 - \mu)}{2} \frac{1 - \delta^2 + \mu(1 + 2\delta + \delta^2)}{1 + \delta + \mu(1 - \delta)} r - c$$

Simplify to:

$$\frac{1 - \mu_L}{1 - \mu} \geq \frac{1 - \delta^2 + \mu(1 + 2\delta + \delta^2)}{1 + \delta + \mu(1 - \delta)} - \frac{2c}{r(1 - \mu)}$$

The condition for a high complexity equilibrium with cost c is written:

$$\frac{1 - \mu_H}{2} r - c \geq \frac{(1 + \delta)(1 - \mu)}{2} r$$

Simplify to:

$$\frac{\mu - \mu_H}{1 - \mu} - \frac{2c}{r(1 - \mu)} \geq \delta$$

Increasing c makes the condition for high complexity more stringent and the condition for low complexity less stringent. Inserting the revised payoffs into the mixed strategy condition outlined in Section 3.4, confirms that an increase in the cost of obfuscation leads firms to obfuscate less. ■

3.9 Appendix B: Alternative Characterisation of Firm Profit.

When characterising obfuscation strategies in Appendix A it is more convenient to write the expected profit of the more complex firm at the reservation price ($p_j \rightarrow r$), instead of the lower bound of the pricing distribution. This follows Narasimhan (1988). The following equivalency condition facilitates this:

$$\frac{(1+\delta)(1-\mu)r}{1+\delta+\mu(1-\delta)} \cdot \left(\frac{1-\delta+\mu+\delta\mu}{2} \right) = \left(\frac{(1-\delta)(1-\mu)}{2} + \underbrace{\mu \left(\frac{2\delta(1-\mu)}{1+\delta+\mu(1-\delta)} \right)}_{\text{Mass Point}} \right) r$$

The L.H.S. is the expected payoff when the firm prices at the lower bound of the price distribution. The R.H.S. is the expected payoff when the firm prices at the upper bound of the price distribution ($p_j \rightarrow r$). Simplify the R.H.S. to yield:

$$\frac{(1+\delta)(1-\mu)}{1+\delta+\mu(1-\delta)} \cdot \left(\frac{1-\delta+\mu+\delta\mu}{2} \right) = \left(\frac{(1-\mu)(1+\mu+2\delta\mu+\delta^2\mu-\delta^2)}{2(1+\delta+\mu(1-\delta))} \right)$$

It can be demonstrated that the expressions are equivalent. Therefore, the payoff for the most complex firm, when firms select different complexity levels, may be written more conveniently in the form:

$$E[\pi_j | k_i < k_j] = \left[\frac{(1-\mu)(1-\delta^2+\mu(1+2\delta+\delta^2))}{2(1+\delta+\mu(1-\delta))} \right] r$$

■

3.10 Appendix C: Supplementary Material

In this supplementary appendix we demonstrate that the parameters governing consumer sophistication are interactive in determining the direction of change of the firms' payoff with respect to changes in the degree of bias. This series of examples highlights that the absolute values and relative values of the parameters are relevant in determining the impact of the simplicity bias on profit, which prevents a simple condition being written to characterise when the bias will lead to higher or lower profit for firms. Write the parameters in the form: (μ_L, μ, μ_H) . We proceed by writing simple claims that are supported by numerical examples.

Claim 1: The absolute value of μ_L is not the sole determinant of the relationship between firm payoff and the degree of bias.

Example: The payoff function is decreasing for (0.9, 0.6, 0.4) but not decreasing for (0.9, 0.7, 0.6). Therefore the absolute value of μ_L is not the sole determinant.

Claim 2: The absolute value of μ is not the sole determinant of the relationship between firm payoff and the degree of bias.

Example: The payoff function is decreasing for (0.75, 0.5, 0.4) but not decreasing for (0.75, 0.5, 0.48).

Claim 3: The absolute value of μ_H is not the sole determinant of the relationship between firm payoff and the degree of bias.

Example: The payoff function is decreasing for (0.8, 0.7, 0.5) but not decreasing for (0.7, 0.6, 0.5).

Claim 4: The difference between μ and μ_H is not the sole determinant of the relationship between firm payoff and the degree of bias.

Example: The payoff function is decreasing for (0.7, 0.4, 0.3) but not decreasing for (0.9, 0.7, 0.6).

Claim 5: The difference between between μ_L and μ is not the sole determinant of the relationship between firm payoff and the degree of bias.

Example: The payoff function is decreasing for (0.9, 0.7, 0.5) but not decreasing for (0.9, 0.7, 0.6).

Summary: In this supplementary analysis we demonstrate that the parameters governing consumer sophistication are interactive in determining the impact of an increase in the degree of bias

on firm profit. This prevents a single algebraic condition from specifying the exact cut-off from which the bias may increase or decrease firm profit.

Chapter 4

Promotions with Biased Consumers

4.1 Introduction

The presentation of product information has a significant influence on consumer decision-making. In particular, firms may withhold relevant price information or present information in a complicated way, to limit consumers' ability to appraise competing offers for homogeneous products (Ellison, 2005; Gabaix and Laibson, 2006). Industrial economists are well versed in the motivations for these practices, referred to as strategic obfuscation. By introducing spurious complexity into their pricing schemes, firms are able to manufacture an environment of increasingly imperfect consumer decision-making, which leads to fewer consumers purchasing from the cheapest seller. Price competition is weakened and firms earn higher profits, which generates a collective incentive for firms to obfuscate.

Common examples of obfuscation include firms framing prices in a complicated way (Ellison and Ellison, 2009; Piccione and Spiegler, 2012; Chioveanu and Zhou, 2013), concealing some dimensions of a product (Gabaix and Laibson, 2006), injecting variance into the prices or quality of multi-attribute products (Spiegler, 2006; Carlin, 2009) and updating promotions frequently (Carlin and Manso, 2011).

A distinct and growing strand of experimental and empirical literature relating to consumer choice behaviour, identifies a bias in consumers' decision-making towards simple products and promotions (Iyengar and Kamenica, 2010; Homburg et al., 2014). In this article we provide a detailed treatment of the interaction between simplicity biased consumers and the obfuscation incentives of firms.¹ In particular, we study the tensions that arise for firms when consumers who cannot understand prices are biased towards simple promotion schemes, beyond what has been considered before.

Specifically, our contribution in this article is to address the following important questions: How does the presence of simplicity biased consumers affect pricing strategies, firm profit and

¹The modelling approach of this chapter differs strongly from the structure presented in Chapter 3. The main distinctions from Chapter 3 are discussed in significant detail in Section 4.1.1.

consumer welfare, when obfuscation is instantly adjustable? Does the presence of simplicity biased consumers lead to lower obfuscation by firms? Are obfuscation incentives stronger when firms choose complexity and prices simultaneously or sequentially?² Can an obfuscation-free market be sustained by market forces in the absence of policy interventions? Do policies to improve consumer sophistication enhance consumer welfare and market transparency?

In an extension to the main model we also consider the opportunities for firms to accompany their product and its price with strategic marketing, which influences consumers' perceptions of the complexity of their product. This facilitates an examination of the interaction between pricing, obfuscation and advertising incentives. This approach differs markedly from previous studies that model advertising and obfuscation techniques as a single parameter for firms, rather than individual choice variables. That is, a firm in our model may create a complicated product that is marketed to appear simple to consumers. This formalises the ideas presented in Chapter 3.

In the model two firms sell a homogeneous product to consumers and simultaneously select both their price and complexity, based on the elegant framework of Chioveanu and Zhou (2013). The complexity of the firms' offers aggregates to market complexity, which determines the proportion of consumers that are able to accurately evaluate prices. Specifically, for every level of market complexity there are two types of consumers: Informed and Uninformed. Informed consumers are able to compare prices and buy from the cheapest seller. Uninformed consumers are impaired by market complexity and cannot compare any of the prices.

Within the literature, uninformed consumers typically buy from a randomly selected firm because sellers are equally prominent and consumers have no brand preferences. Buyers also cannot infer the relative price from the relative complexity of a firm's promotion. However, when firms are able to obfuscate, the random purchase rule implies an additional assumption: Uninformed consumers are indifferent between simple and complex products. This conflicts with experimental and empirical evidence. Therefore, we rescind this assumption and allow uninformed consumers to be biased towards less complex promotions.

Homburg et al. (2014) explore the role of price complexity on consumer decision-making using experiments. They find that individuals tend to overestimate the price of complicated options and underestimate the price of simple options. Therefore, buyers associate more complex pricing schedules with higher prices, lower price fairness and lower firm transparency, which creates a bias towards simple promotions. Iyengar and Kamenica (2010) also show that individuals are biased towards intrinsically less complex investment products when selecting between retirement investment funds. In addition, this preference for simplicity strengthens as the decision becomes more complex due to an increase in choice. De los Santos et al. (2012) also find that the ease with which consumers can navigate around prices on online interfaces plays a significant role in

²Equivalently, do firms obfuscate more when obfuscation is quickly adjustable or a long term commitment?

decision-making. Huck and Wallace (2015) also show that price framing has a significant influence on individuals' decision-making in experiments.

There exists many inter-related explanations for simplicity biases, which are likely to vary across markets. These include: an intrinsic preference for simple products, additional utility from obtaining a simple product, ambiguity aversion and biases in consumers' estimation calculations. However, the model we present is intentionally general and is not predicated on any particular origin of the bias.

The increasing importance of simplicity as a product differentiation tool for firms is also evident from the surge in simplicity orientated marketing campaigns. Ryanair, a budget airline, advertises "Low Fares, Made Simple."³ Santander, a high street bank, advertises "Savings made simple."⁴ Energy provider NPower claim "We're making our tariffs simpler."⁵ Groceries firm, Sainsbury's, also relaunched their own-brand products as a 'simplicity' range, removed jargon from their promotions and created a simplicity-centric telecommunications brand.⁶ In its core grocery arm, Sainsbury's is also phasing out multi-buy discounts in favour of simple fixed prices,⁷ with rivals Tesco and Asda following closely behind in an attempt to regain customers from discounters.⁸ The growing demand for simplicity also provides new opportunities for innovation. 'SellSimpleEstateAgency.com' and 'HouseSimple.com' now provide a cheaper and less complex platform for homeowner's to sell residential property.

Using the model, we show that the introduction of simplicity biased consumers leads firms to compete in their complexity levels, in addition to prices and reduces firms' ability to obfuscate. This leads to more consumers being able to understand prices. As a consequence, firms compete more fiercely in the price dimension at the expense of profit. In equilibrium, firms combine high prices with high complexity because a firm that is less competitively priced benefits the most from consumer confusion. This implies that the simplicity bias is rational for confused consumers because less complex products are cheaper. Therefore, biased consumers obtain a lower price than their random purchase counterparts and generate a multiplied increase in consumer welfare because: (i) they buy from the cheapest seller, (ii) they create a positive market externality by restricting the obfuscation of firms, which allows more consumers to understand prices and further stimulates price competition.

³ Available at: m.thedrum.com/news/2015/02/04/ryanair-s-pan-european-creative-brief-goes-dare TSB. Accessed 04/05/16.

⁴ In-store promotional materials. Accessed 31/03/16.

⁵ www.npower.com/tariff-simplification/why-have-we-done-this/index.htm. Accessed 04/05/16.

⁶ www.j-sainsbury.co.uk/media/latest-stories/2013/20130701-sainsburys-launches-mobile-by-sainsburys. Accessed 04/05/16. www.kansomarketing.com/blog/sainsburys-value-simplicity. Accessed 04/05/16.

⁷ www.j-sainsbury.co.uk/media/latest-stories/2016/0211-sainsburys-to-phase-out-multi-buy-promotions-in-favour-of-lower-regular-prices/. Accessed 05/05/16.

⁸ www.telegraph.co.uk/finance/personalfinance/household-bills/12151067/Supermarkets-face-ban-on-special-offers-that-cost-shoppers-1000-a-year.html. Accessed 05/05/16. www.ibtimes.co.uk/tesco-plans-simple-pricing-win-back-shoppers-1541131. Accessed 05/05/16.

In the extension, we consider two marketing opportunities that arise when consumers exhibit a simplicity bias. First, we allow firms to influence the degree of bias exhibited by uninformed consumers using long-term marketing campaigns. This simple extension discerns an additional layer of incentives in strategic marketing: Firms choose to offer a simple product to appeal to consumers' preferences for simplicity, but choose not to encourage consumers towards simple products, which allows firms to obfuscate more and charge higher prices. This line of thought can explain why firms offer a simple product but do not directly encourage consumers to base their decision-making on simplicity. Second, we allow firms to advertise the simplicity of their product, in addition to their price, using a wider marketing campaign. This divorces obfuscation strategies from advertising. We show that this leads to a non-standard element of 'false' advertising whereby sellers couple complex pricing structures with an advert for product simplicity. This allows firms to simultaneously attract biased buyers and limit consumer sophistication.

These results are relevant for policy because obfuscation is a growing concern for regulators. The Competition & Markets Authority (UK) reports that consumer switching is particularly low across current accounts and cite complex pricing and limited information as the dominant cause.⁹ This leads to a weak competitive environment and exacerbates consumer disengagement. Similar practices are observed across the domestic energy market,¹⁰ broadband service industry and notably the groceries sector, which culminated in a supercomplaint from consumer body 'Which?' in 2015.¹¹ Furthermore, existing remedies to obfuscation by retailers tend to induce unintended adverse consequences (Armstrong et al., 2009; Gu and Wenzel, 2014). This places a renewed focus on understanding firms' incentives to develop incentive-compatible competition and consumer protection policies.

Moreover, in this article we emphasise three policy considerations. Firstly, policies to remove behavioural biases in decision-making do not necessarily benefit consumers. In our setting such policies lead to higher obfuscation, higher prices and lower welfare for consumers. Secondly, policies to improve consumer sophistication may induce firms to increase their obfuscation. Therefore, the introduction of price comparison websites and decision-making aids may be less effective than anticipated. This highlights the importance of considering the best response of sellers to market interventions. Our explanation for firms to increase their obfuscation when consumers become more sophisticated also differs from previous accounts predicated on *ex ante* heterogeneous firms (Gu and Wenzel, 2014). Thirdly, firms may engineer complex products that are marketed as simple options for consumers, which constitutes a form of false advertising.

⁹www.gov.uk/government/news/cma-proposes-better-deal-for-bank-customers. Accessed 03/16

¹⁰www.gov.uk/government/news/cma-sets-out-case-for-energy-market-reform. Accessed 03/16

¹¹www.gov.uk/cma-cases/groceries-pricing-super-complaint. Accessed 03/16

4.1.1 Related Literature

In a related article (presented in Chapter 3), we study a variant of the model in which firms must commit to a complexity level before competing in prices. This captures dimensions of obfuscation that cannot be adjusted as quickly as prices and might include the complexity of a website or the terms and conditions presented with the price. In contrast, this fully simultaneous framework captures dimensions of obfuscation that can be adjusted as quickly as the price. This might include the quantity of dimensions in a multi-part tariff.

In the sequential framework presented in Chapter 3, simple products are the most expensive. The reason is that when a firm chooses its price in the second stage, the firm with the lowest complexity is guaranteed a larger share of the uninformed buyers who purchase at any equilibrium price. Therefore, competing for informed consumers is more costly for the least complex firm in terms of the foregone revenue from guaranteed sales, generating a negative correlation between prices and obfuscation at the firm level. As a consequence, consumers who select the simplest product in that framework will buy at the highest average price.

The impact of simplicity seeking consumers on firm profit in the sequential model in Chapter 3, however, is more involved and depends on the outcome of two competing forces. In one direction, simplicity seeking individuals pay a higher price than if they purchased randomly, which increases firm profit (Direct effect). However, simultaneously these buyers exert a positive externality on the market by tempering firms' incentives to obfuscate (Externality effect). This leads to more consumers acting as informed buyers, which creates more fierce price competition and reduces profit.

In the model we study in this chapter, the direct and externality forces operate unidirectionally in favour of the consumer. Therefore, firms always earn lower profits when consumers are biased towards simple products. We show that this leads to very different obfuscation equilibria and advertising behaviour when firms are able to tamper with consumers' perceptions of the complexity of their product using marketing.

This study relates to the growing literature on the role of framing effects in competition (Eliaz and Spiegler, 2011a; Piccione and Spiegler, 2012; Spiegler, 2014). Most closely related is Chioveanu and Zhou (2013) who provide a comprehensive treatment of the influence of various sources of consumer confusion on the optimal price framing strategies of firms. In particular, our model collapses to their duopoly framework when consumers are confused by complex price frames but consumers exhibit no bias towards simple products. However, by focusing on the role of simplicity biased consumers we identify novel dimensions of competition and demonstrate that the incentive for firms to obfuscate may be significantly weaker than anticipated. In an oligopoly extension, Chioveanu and Zhou (2013) also allow uninformed consumers to favour firms with a simple price frame but they do not analyse the impact of the bias on the strategies and payoffs of

firms, which is the primary focus of this study. Furthermore, our model endogenises the degree of simplicity bias as a function of firms' marketing efforts to explore the impact of simplicity biases on firms' incentives in further detail.

Most importantly, models of price-frame competition also treat advertising and framing as a single parameter for firms. In this paper, we distinguish between framing effects and strategic advertising and allow these parameters to exist simultaneously. This generates sharp and empirically testable predictions relating to false advertising.

More broadly, this article contributes to the wider class of studies that analyse obfuscation by firms. Ellison (2005) shows that firms may lure consumers using low quality products before offering high priced upgrades at the point of purchase. Gabaix and Laibson (2006) show that firms will compete intensely in the market for base products but conceal the existence of high priced add-ons from consumers until consumers are committed to the base good. Spiegel (2006) also explains why firms may inject variance into the quality of multi-part products when consumers make decisions based on only a limited subset of the products' characteristics. More recently, Carlin (2009) describes how firms may add complexity to their pricing structures to soften competition for financial products. The main distinction with this literature is that in these frameworks there exists no role for confused consumers to prefer simple products.

Several studies demonstrate that the incentive for firms to obfuscate add-on products may be weakened in a Gabaix and Laibson (2006) market with add-ons but the explanations differ from ours. In Dahremoller (2013), firms vary in their profitability in the add-on market, which incentivises one firm to inform consumers regarding the add-on. In Wenzel (2014), the proportion of uninformed buyers that become aware of the add-on is an increasing function of the number of unshrouding firms, rather than a fixed proportion, which admits a wider range of unshrouding equilibria. In our model, however, firms are *ex ante* symmetric and only one product is sold. Firms choose not to obfuscate because of competition in the obfuscation dimension, created by the simplicity bias. In the search literature, Wilson (2010) and Ellison and Wolitzky (2012) also study models in which firms obfuscate their product by inflating consumers' search costs.

Carlin and Manso (2011) provide an interesting explanation for firms to frequently change their promotions, which differs from ours. In their model, consumers are initially divided into informed and uninformed types but over time more consumers become informed due to learning effects. Firms respond by changing their promotions to reset the consumer learning process and soften competition. In our model, consumer learning is not directly captured and firms change both their price and the complexity of their promotion to compete for biased consumers. However, the idea that uninformed consumers begin to infer that more complex promotions are more expensive, in the sense of learning over time, is consistent with the equilibrium we characterise.¹²

¹²This argument is further developed in Section 4.5.

The idea that firms may choose to simplify their products also develops our understanding of the adoption of common standards by retailers. Gaudeul and Sugden (2012) explore a model in which firms obfuscate by deviating to an individuated promotion format that is more complicated for consumers to compare. They show that consumers are rational in discounting firms that differentiate their promotions because these firms charge higher prices in equilibrium. However, there are three main differences with the present study.

First, obfuscation adopts a broader role in our model, as any complexity that sellers are able to introduce into their product or price, in addition to individuated promotions. Second, the model considers a role for firms to influence the degree of bias exhibited by uninformed consumers using competitive marketing. Third, this paper outlines the impact of measures to improve consumer sophistication on the profit and strategies of firms. In particular, our model demonstrates that firms will offer simple products but discourage consumers from buying from the least complex firm.

Simplicity may also be interpreted as a method for firms to increase the saliency of their product to buyers, which links to the advertising literature (Haan and Moraga-González, 2011). In this direction, simplicity plays a product differentiation role by appealing to uninformed buyers. However, simplicity in our model adopts a secondary effect beyond product differentiation by influencing the proportion of informed and uninformed buyers.

This interpretation of simplicity also connects to the literature on obfuscation amongst firms that vary *ex ante* in prominence. Gu and Wenzel (2014) show that when firms select an obfuscation level before setting prices and one firm is guaranteed a larger share of uninformed buyers, independently of its obfuscation, interpreted as prominence, the more prominent seller selects the highest obfuscation and price. In our model, however, prominence is determined endogenously according to the relative obfuscation levels of firms and the more prominent firm charges a lower price.

The simplicity bias in our model also relates to recent models in which consumers behave as ‘salient-thinkers’ (Bordalo et al., 2013). In this literature, consumers pay greater attention to dimensions of a product that differ most from the average value of that dimension in the market. In our framework, simplicity becomes attractive to consumers whenever firms add complexity into their prices, causing sellers to differ in their complexity. If all firms adopted the same constant obfuscation level, the simplicity bias would play no direct role because all firms are equally complex and confused consumers return to buying randomly. However, the bias may remain active indirectly, by disciplining firms to offer this fixed obfuscation level. We show that this occurs when firms are constrained to selecting the minimum level of obfuscation, which arises when the bias is sufficiently large.

Most broadly this article fits within the behavioural industrial organisation literature in which sophisticated firms engage with biased and imperfectly informed consumers (Spiegler, 2011). Furthermore, the article connects to the topic of exploitative contracts studied in behavioural contract theory. See Köszegi (2014) for a recent survey.

The article is structured as follows. Section 4.2 describes the main model and characterises the main equilibria. Section 4.3 analyses the comparative statics and in particular the role of the simplicity bias on firm conduct and consumer welfare. Section 4.4 considers an extension in which firms are able to influence the degree of simplicity bias using marketing strategies. The results are discussed in Section 4.5 and Section 4.6 concludes.

4.2 The Model

Consider a market with two rational, risk neutral, profit maximising firms indexed $i, j \in N$, who sell a single homogeneous product to a unit mass population of consumers. Sellers are equally prominent and are able to satisfy all of the demand they acquire. The cost of production is also symmetric and without loss of generality, normalised to zero. The strategic variables for each firm $i \in N$ are a price-complexity pair (p_i, k_i) . The choice of complexity is also costless and for tractability firms make a binary choice between high (H) and low (L) complexity.

Consumers will each purchase exactly one unit of the product at any price not exceeding their common reservation threshold ($r = 1$) and they are not perturbed from purchasing by the complexity that they may face in the market. This fits well with many retail environments, where consumers are sufficiently compelled to make a purchase providing the price is below what they are willing to pay. This includes mobile phone contracts, domestic energy services and financial products.

Consumers do not exhibit any brand preferences and wish to purchase from the lowest priced seller but for a given level of market complexity, which is strictly increasing in the complexity of each firm, only some are able to perfectly evaluate prices.¹³ Informed consumers, $\mu_{(k_i, k_j)} \in (0, 1)$, identify and purchase from the firm with the lowest price but uninformed consumers, $(1 - \mu_{(k_i, k_j)})$, are impaired by market complexity and cannot compare any prices.¹⁴

Diverging from mainstream studies, we dispense with the assumption that uninformed consumers purchase randomly. Instead, some of the uninformed consumers: $\delta(1 - \mu)$, $\delta \in [0, 1]$, purchase from the firm with the lowest complexity. We term these consumers simplicity biased. The remaining group: $(1 - \delta)(1 - \mu)$, are distributed evenly across the firms such that the random purchase rule is restored whenever: (i) $\delta = 0$ or (ii) all firms select the same complexity level. The

¹³De los Santos et al. (2012) provide evidence that consumers follow a fixed sample search process, not a sequential updating search style. This validates the 'all-or-nothing' approach.

¹⁴For notation, the dependence of μ on k_i and k_j will often be suppressed.

division of uninformed consumers between simplicity biased and random purchasers is given exogenously according to a fixed proportional rule (δ) in the main model and a function of firms marketing efforts in an extension in Section 4.4.

The assumption that uninformed consumers cannot infer the relative price of a firm from their relative complexity is reasonable in many circumstances. However, it emerges in equilibrium that a less complex firm actually charges a lower price than a more complex firm. Therefore, the simplicity bias could also be explained as consumers who repeatedly engage with the market inferring which firm will offer the lower price based on observable complexity levels, which introduces an element of consumer learning (Carlin and Manso, 2011). The assumption that uninformed consumers can distinguish between complicated and simple options is also plausible and not unique to our study (Wilson, 2010).

Consumers' utility is an inverse function of prices, measured by the difference between their reservation threshold and the price they pay. In practice, consumers may derive additional utility from the consumption of a simple product. In Section 4.5 we show that this perturbation does not qualitatively affect our results. In addition, if consumers disagree which firm offers the most simple product, this is equivalent to some uninformed consumers preferring the complex option. The main results continue to hold with a reduced degree of bias, providing more uninformed consumers prefer the low complexity firm. The intuition is that the disagreeing uninformed buyers cancel out in terms of the degree of bias.

Write the complexity decisions of each firm as the pair $\{k_i, k_j\}$, such that there exists three possible levels of market complexity: $\{L, L\}$, $\{L, H\} = \{H, L\}$, $\{H, H\}$. Each level of market complexity maps to a corresponding value of informed and uninformed consumers. Define the proportion of informed consumers as μ_L if both firms adopt low complexity, μ if only one firm adopts low complexity and μ_H if both firms adopt high complexity where: $0 < \mu_H < \mu < \mu_L < 1$. This directly captures the reduction in consumers' ability to compare prices when sellers complicate their pricing. In addition, this distinguishes the present framework from models in which firms select a price and a unit of measurement for the price, and consumers are only impaired when firms frame their prices differently such that $\mu < \mu_H = \mu_L$, as described by Chioveanu and Zhou (2013).

For tractability we initially assume that consumers are impaired by obfuscation at a constant rate such that $\mu_L - \mu = \mu - \mu_H$, which is not unique to our study (e.g. Gu and Wenzel (2014)). This condition simplifies the characterisation procedure and elucidates the influence of the simplicity bias more clearly, which is the principal focus of this study. In Section 4.2.5 we relax this assumption. The following timeline summarises the dynamics of the symmetric simultaneous-move two-period game:

Period 1:

- Each firm $i \in N = \{1, 2\}$ selects a price p_i and a complexity level $k_i \in \{L, H\}$ for their product.¹⁵
- The proportion of informed consumers is determined as a declining function of market complexity such that: $0 < \mu_H < \mu < \mu_L < 1$, where $\{L, L\} = \mu_L$, $\{L, H\} = \{H, L\} = \mu$, $\{H, H\} = \mu_H$.

Period 2:

- Consumers purchase according to their type:
 - $\mu_{(k_i, k_j)} \in (0, 1)$ are Informed and buy from the firm with the lowest price.
 - $\delta(1 - \mu_{(k_i, k_j)})$ are Simplicity Biased and buy from the firm with the lowest complexity.
 - $(1 - \delta)(1 - \mu_{(k_i, k_j)})$ are Random Purchasers and buy from a randomly selected firm.
- Firms receive their respective profit.

The equilibrium depends critically on the degree of simplicity bias exhibited by uninformed consumers. When the degree of bias is low, firms engage in a mixed strategy over complexity with prices drawn from marginal probability functions that are specific to high and low complexity (Proposition 1). As the strength of the simplicity bias increases, the probability that each firm chooses low complexity increases until the bias reaches the critical value δ^c . At δ^c the mixed strategy distribution becomes degenerate and all firms select low complexity. Firms continue to price stochastically according to a marginal distribution function to prevent the rival from undercutting by a marginal deviation. For a larger bias, $\delta > \delta^c$, firms continue to select low complexity as a pure strategy Nash equilibrium (Proposition 2). Let $Supp F(p)$ denote the support of the cumulative probability function $F(p)$ that allocates probability across the optimal interval of prices.

Proposition 1. Define $\delta^c \equiv \frac{\mu - \mu_H}{1 - \mu}$. For $\delta < \delta^c$ there exists a mixed strategy equilibrium in which firms select low complexity with probability $\lambda_L = \frac{\mu - \mu_H + \delta - \delta\mu}{\mu_L - \mu_H} \in (0, 1)$ and high complexity with probability $\lambda_H = 1 - \lambda_L$. When low complexity is selected, a price p' is drawn randomly from a distribution $F^L(p')$. When high complexity is selected, a price p is drawn randomly from a distribution $F^H(p)$. The price distributions have adjacent supports such that $Supp F^L(p') = [p^l, p^m]$, $Supp F^H(p) = [p^m, r]$.

¹⁵For consistency we refer to the complexity of the product, which with some interpretations of obfuscation may equivalently be stated as the complexity of the pricing scheme.

The supports of the price distributions are illustrated in Figure 4.1.




Figure 4.1: Equilibrium Pricing: $\delta < \delta^c$

The characterisation proceeds in steps following Chioveanu and Zhou (2013). First we write the equilibrium payoff to each firm at the price-complexity pair $(H, 1)$, which is demonstrated to be an element of the mixed strategy equilibrium. Second, we compute the expected payoff for every other price-complexity pair (H, p) and (L, p') , where $p \in \text{Supp } F^H(p)$ and $p' \in \text{Supp } F^L(p')$. By definition of a mixed strategy the expected payoff at every price-complexity pair in the equilibrium must be equal, which produces a set of indifference conditions.

Third, we derive explicit expressions for p^l and p^m where p^l is the lower bound of $F^L(p')$ and p^m is the upper (lower) bound of $F^L(p')$ ($F^H(p)$). We also compute the optimal probability with which each firm will select low complexity, λ_L . Fourth, we demonstrate the absence of a profitable deviation to a strategy in which a complexity level is combined with a price from the support of the other distribution and determine an explicit expression for profit in terms of μ_L , μ , μ_H and δ . Finally, closed form solutions for the marginal price distributions can be computed.

4.2.1 Equilibrium Payoffs

A strategy for each firm comprises three elements: λ_L indicates the probability assigned to low complexity, $F^L(p')$ indicates the probability distribution over prices associated with low complexity and $F^H(p)$ indicates the probability distribution over prices associated with high complexity. By standard arguments, there can exist no pure strategy equilibrium in prices. Pricing deterministically exposes the firm to the usual marginal undercutting deviation by a rival, leading to a mixed strategy solution. The following lemma provides the foundation for the characterisation. All proofs are contained in the appendix.

Lemma 1. *For $\delta < \delta^c$, the price-complexity pair $(H, 1)$ belongs to the equilibrium mixed strategy.*

The intuition is that in a symmetric equilibrium with a positive fraction of uninformed consumers, the maximum price that any firm will set is the reservation price, $r = 1$. At any price above r , consumers do not purchase from the firm. At any common maximum price (\bar{p}) below 1, the firm's demand consists exclusively of uninformed buyers because the firm charges the highest price with certainty. These individuals would also purchase at any price not exceeding r . Therefore,

increasing the price to r yields no reduction in the quantity of sales and a strict increase in profit, contradicting \bar{p} as the upper bound.

The reason that a firm will combine the maximum price with high complexity is that when a firm charges the maximum price it never attracts informed consumers. Therefore by selecting high complexity, the value of informed consumers is minimised, and the firms' potential client base (uninformed consumers) is maximised. In contrast, a firm that offers a lower price is competing directly for informed consumers because this firm requires the increase in demand to sufficiently compensate for the reduction in profit per sale, leading the firm to select a lower level of complexity.

For the strategy $(H, 1)$, the firm sells only to uninformed consumers but the distribution of uninformed consumers between the firms depends on the complexity choice of the rival. When the competitor selects high complexity, the demand of uninformed consumers is shared equally amongst the two firms. When the competitor selects low complexity, the rival attracts all simplicity biased consumers and the firm secures only the demand from random purchaser's. Let λ_L be the probability that firm j selects low complexity in equilibrium. The expected payoff for firm i , denoted $\pi_i(k_i, p_i)$ is given by:

$$\pi_i(H, 1) = \frac{1}{2} \cdot [\lambda_L \cdot (1 - \delta)(1 - \mu) + \lambda_H \cdot (1 - \mu_H)] \quad (4.1)$$

When a firm adopts high complexity, the mixed strategy equilibrium must ensure that the firm is indifferent between charging 1 and any other price p contained in the support of the marginal pricing distribution, $F^H(p)$. At any price p , the firm continues to attract the same share of uninformed consumers, contingent on the behaviour of the rival. However, the firm has an additional revenue stream from the potential sales to informed buyers. Specifically, if the rival selects high complexity, the firm offers the lowest price with probability $[1 - F^H(p)]$, which yields the following expected payoff:¹⁶

$$\pi_i(H, p) = \frac{1}{2} \cdot [\lambda_L \cdot (1 - \delta)(1 - \mu) + \lambda_H \cdot (1 + \mu_H [1 - 2F^H(p)])] \cdot p \quad (4.2)$$

Equally, the firm must be indifferent between selecting the pairs $(H, 1)$ and (L, p') , where p' is any price contained in the support of the equilibrium price distribution, $F^L(p')$. The payoff associated with the strategy (L, p') is given by:

$$\pi_i(L, p') = \frac{1}{2} \cdot [\lambda_L \cdot (1 + \mu_L [1 - 2F^L(p')]) + \lambda_H \cdot (1 + \delta + \mu - \delta\mu)] \cdot p' \quad (4.3)$$

¹⁶Supplementary materials detailing the derivation of the payoffs is provided in Appendix B (Section 4.8).

If the rival selects low complexity, uninformed consumers are distributed evenly and firms compete in prices for informed consumers. If the rival selects high complexity, the firm secures an equal share of the consumers who purchase randomly and the entire population of simplicity biased consumers. The firm also secures the custom of all informed consumers with certainty because the more complex rival always charges a higher price. This follows from the adjacent supports of the price distributions. Therefore, by definition of a mixed strategy the following two indifference conditions must hold with equality:

$$\lambda_L \cdot (1 - \delta)(1 - \mu) + \lambda_H \cdot (1 - \mu_H) = [\lambda_L \cdot (1 - \delta)(1 - \mu) + \lambda_H \cdot (1 + \mu_H [1 - 2F^H(p)])] \cdot p \quad (4.4)$$

$$\lambda_L \cdot (1 - \delta)(1 - \mu) + \lambda_H \cdot (1 - \mu_H) = [\lambda_L \cdot (1 + \mu_L [1 - 2F^L(p')]) + \lambda_H \cdot (1 + \delta + \mu - \delta\mu)] \cdot p' \quad (4.5)$$

The first indifference condition makes clear that when a firm reduces its price below r , in an attempt to secure informed consumers, the expected increase in demand must exactly compensate for the reduction in profit per customer, such that the total effect on profit is neutral. The second indifference condition illustrates that a firm that deviates to low complexity will charge a lower price than the high complexity rival, to attract the now-larger proportion of informed consumers.

Solving (4.5) determines explicit solutions for p^l and p^m , where $F^H(p^m) = 0$, $F^L(p^m) = 1$ and $F^L(p^l) = 0$.

$$p^m = \frac{\lambda_L \cdot (1 - \delta)(1 - \mu) + \lambda_H \cdot (1 - \mu_H)}{\lambda_L \cdot (1 - \mu_L) + \lambda_H \cdot (1 + \delta + \mu - \delta\mu)} \quad (4.6)$$

$$p^l = \frac{\lambda_L \cdot (1 - \delta)(1 - \mu) + \lambda_H \cdot (1 - \mu_H)}{\lambda_L \cdot (1 + \mu_L) + \lambda_H \cdot (1 + \delta + \mu - \delta\mu)} \quad (4.7)$$

It is clear that $p^m > p^l$, which is consistent with the equilibrium proposition. In addition, when $\lambda_L = 1$, $p^m = 1$, which is made clearer when p^m is derived from (4.4):

$$p^m = \frac{\lambda_L \cdot (1 - \delta)(1 - \mu) + \lambda_H \cdot (1 - \mu_H)}{\lambda_L \cdot (1 - \delta)(1 - \mu) + \lambda_H \cdot (1 + \mu_H)} \quad (4.8)$$

It can be shown that the two expressions for the intermediate price bound, p^m , are equivalent.¹⁷

¹⁷The derivation is contained in Appendix B (Section 4.8).

4.2.2 Equilibrium Obfuscation

It is now possible to derive an expression for λ_L and demonstrate that a firm will never combine a price $p \in \text{Supp } F^H(p)$ with low complexity and correspondingly will never combine a price $p' \in \text{Supp } F^L(p')$ with high complexity. This translates into two inequalities: $\pi_i(H, p) \geq \pi_i(L, p)$ and $\pi_i(L, p') \geq \pi_i(H, p')$, which must be satisfied simultaneously. Inserting the relevant expressions yields:

$$\lambda_L \cdot (1 - \delta)(1 - \mu) + \lambda_H \cdot (1 + \mu_H [1 - 2F^H(p)]) \geq \lambda_L \cdot (1 - \mu_L) + \lambda_H \cdot (1 + \delta + \mu [1 - \delta - 2F^H(p)]) \quad (4.9)$$

$$\lambda_L \cdot (1 + \mu_L [1 - 2F^L(p')]) + \lambda_H \cdot (1 + \delta + \mu - \delta\mu) \geq \lambda_L \cdot (1 - \delta + \mu [1 + \delta - 2F^L(p')]) + \lambda_H \cdot (1 + \mu_H) \quad (4.10)$$

The supports of the price distributions are adjacent, which implies that the firm must be indifferent between low and high complexity at price p^m . Therefore, the above two conditions must hold with equality at p^m where $F^H(p^m) = 0$ and $F^L(p^m) = 1$. The intuition is that the firm is indifferent between high and low complexity when the boundary price of the distributions is selected. Inserting these results into either expression yields an exact solution to λ_L :

$$\lambda_L = \frac{\mu - \mu_H + \delta - \delta\mu}{\mu_L - \mu_H} \quad (4.11)$$

λ_L is a probability and must be restricted to the closed interval $[0, 1]$, which is satisfied for any $\delta \leq \delta^c$. The observation that $\lambda_L = 1$ whenever $\delta = \delta^c = \frac{\mu - \mu_H}{1 - \mu}$, is intuitive, once rewritten in the form: $\delta(1 - \mu) = \mu - \mu_H$. The left hand side corresponds to the value of simplicity biased consumers and the right hand side corresponds to the marginal effect of obfuscation by a firm. When the benefit from obfuscating in terms of the reduction in consumer sophistication, is exactly equal to the cost of obfuscation, in terms of the foregone demand of simplicity biased consumers, the incentive for a firm to select high complexity disappears.

The expression for λ_L in (4.11) can be inserted directly into (4.9) and (4.10) to confirm the optimality of the obfuscation-price pairs at every price in their respective supports. The proof is contained in the Appendix.

4.2.3 Derestricting the Simplicity Bias.

Proposition 2. For $\delta \geq \delta^c$ there exists a pure strategy equilibrium in obfuscation in which all firms select low complexity and prices are drawn stochastically from the cumulative probability function: $F^L(p')$, with continuous and atomless support: $\left[\frac{1-\mu_L}{1+\mu_L}, r\right]$.

Proposition 2 implies that a further increase in the degree of simplicity bias beyond the level required for firms to exclusively select low complexity, never induces firms to return to a strategy that includes high complexity. That is, when more consumers prefer a simple product, it is not profitable for a firm to increase their obfuscation level. Moreover, increasing δ beyond δ^c has no impact on the pricing behaviour or profit of sellers. When $\lambda_L = 1, p^m = 1$ and the price distribution collapses to:

$$F^L(p') = 1 - \frac{1 - \mu_L}{2\mu_L} \left[\frac{1 - p'}{p'} \right] \quad (4.12)$$

with support $\left[\frac{1-\mu_L}{1+\mu_L}, r\right]$.

In Section 4.3, we refer to this outcome as a low complexity equilibrium. This equilibrium also resembles the pricing behaviour in Varian's (1980) classic model of sales. In that model, firms maximise profit by choosing prices with a fixed proportion of informed and uninformed consumers. The key difference is that the level of consumer sophistication is determined endogenously in our framework. For the remainder of this article we focus primarily on the more interesting case of $\delta < \delta^c$. For generality define λ_L in the form:

$$\lambda_L = \text{Min} \left\{ \frac{\mu - \mu_H + \delta - \delta\mu}{\mu_L - \mu_H}, 1 \right\} \quad \lambda_H = \text{Max} \left\{ \frac{\mu_L - \mu + \delta\mu - \delta}{\mu_L - \mu_H}, 0 \right\} \quad (4.13)$$

4.2.4 Firm Profit & Price Distributions

We are now able to insert the expression for λ_L into (4.1) to derive a closed form solution for the expected payoff to each firm:

$$\pi_i = \frac{1}{2} \cdot \left[1 - \mu_H - \frac{(\mu - \mu_H + \delta - \delta\mu)^2}{\mu_L - \mu_H} \right] \quad (4.14)$$

If all consumers are informed, irrespective of firms' obfuscation efforts, firms cannot charge a price above marginal cost. Conversely, if all consumers are uninformed, irrespective of firms' obfusca-

tion behaviour, firms extract the maximum payoff by charging every consumer their reservation price.

Closed form solutions for the marginal distribution functions that assign probability across every price in the support, can also be derived by inserting the explicit solution for λ_L into (4.4) and (4.5):

$$F^H(p) = 1 - \left[\frac{1-p}{2 \cdot p \cdot \mu_H} \right] \left[1 - \mu_H + \frac{(\mu - \mu_H + \delta - \delta\mu)}{\mu_L - \mu - \delta + \delta\mu} (1 - \delta)(1 - \mu) \right] \quad (4.15)$$

defined over the continuous and atomless support $[p^m, 1]$.

$$F^L(p') = \frac{1}{2} \left[1 + \frac{1}{\mu_L} \left(1 - \frac{1}{p'} \cdot \left[(1 - \delta)(1 - \mu) + \frac{\mu_L - \mu + \delta\mu - \delta}{\mu - \mu_H + \delta - \delta\mu} (1 - \mu_H - p' [1 + \delta + \mu - \delta\mu]) \right] \right) \right] \quad (4.16)$$

defined over the continuous and atomless support $[p^l, p^m]$.

It can be verified that $F^H(1) = 1$, $F^H(p^m) = 0$, $F^L(p^m) = 1$, $F^L(p^l) = 0$.¹⁸

4.2.5 Decreasing Marginal Effect of Obfuscation

One can argue that as firms confuse customers using complicated pricing or other forms of obfuscation, increasing the complexity of their products further will generate a smaller additional increase in consumer confusion. Formally, the marginal effect of obfuscation on the quantity of informed buyers may be decreasing at a decreasing rate. We therefore show that the main equilibria are not particularly sensitive to the assumption of a constant effect of obfuscation.

Specifically suppose $\mu_L - \mu > \mu - \mu_H$ such that the marginal effect of introducing spurious complexity is decreasing. Assume also that $\delta < \delta^c$ to focus our attention on the most interesting obfuscation behaviour. It can be demonstrated (See Appendix B) that firms continue to follow the obfuscation and pricing strategy outlined in the main model.

¹⁸This is also clear from the example in Section 4.3. The derivations are given in the Appendix.

4.3 Comparative Statics & Policy Measures

This section explores the comparative statics of the model and their implications for competition and consumer protection policies.

4.3.1 Comparative Statics

Proposition 3. The comparative statics of the simplicity bias are as follows:

1. An increase in the degree of bias leads to:
 - (i) An increase in the probability that firms assign to low complexity.
 - (ii) A decrease in firm profit and an increase in consumer welfare.
 - (iii) An increase in price dispersion: $r - p^l$.
2. Firms are disciplined to a low complexity equilibrium for a smaller bias when prices and obfuscation are chosen simultaneously, than when obfuscation is committed in a preceding stage.
3. For $\delta = 0$, firms earn higher profit when obfuscation is committed before prices, than when obfuscation and prices are selected simultaneously.
4. In a low complexity equilibrium, the profit and pricing strategies of firms are independent of the simultaneous and sequential timing structures.

An increase in the degree of simplicity bias strengthens the incentive for a firm to undercut the complexity level of the rival because more consumers buy from the least complex seller. Therefore, the purchasing behaviour of biased individuals exerts a positive externality on the market by restricting firms' propensity to obfuscate. The reduction in market complexity allows more consumers to compare prices, which stimulates price competition amongst sellers. Prices and profit subsequently decline, which leads to higher welfare for consumers in monetary terms. Furthermore, if consumers accrue additional utility from consuming a simple product, the effect on consumer welfare is amplified.

An increase in the bias also incentivises a firm that selects low complexity to engage in more significant price reductions, to compete for the larger quantity of informed buyers. This causes p^l to decrease. Prices are maximally dispersed when the bias is sufficiently large to discipline firms to low complexity. In this case the softening of price competition when a firm selects high complexity is insufficient to compensate for the loss of demand from biased consumers.

In this simultaneous framework, firms select low complexity with positive probability irrespective of the degree of bias. If a firm is more competitively priced, he prefers that more consumers are able to evaluate prices and buy from the cheapest supplier. If a firm selects a high price, he prefers that more consumers are confused, to conceal his lack of competitiveness.

This differs markedly from the related article in Chapter 3, in which obfuscation is committed before firms compete in prices. In that model, firms do not vary between high and low complexity in the absence of simplicity biased consumers. Instead, firms hold a common interest in confusing as many consumers as possible in the first stage, before competing in prices in the second stage. Therefore, there is no reason to deviate from high complexity in the sequential setting when the simplicity bias is sufficiently small.

The sequential game in Chapter 3 also features a standard prisoner's dilemma: When the bias is sufficiently large to instil a low complexity equilibrium, firms earn higher profits by coordinating to high complexity but the existence of a profitable deviation induces a low complexity equilibrium. This structure is not present in the current simultaneous game.

In both models, an increase in the bias causes firms to select low complexity with greater probability but firms always assign higher weight to low complexity in the simultaneous framework because of the additional pricing-related incentive. This is further reflected in the higher payoff that firms earn in the sequential model than the simultaneous model when the degree of bias is zero. Interestingly, however, when the bias is sufficiently strong and the parameters allow for a low complexity equilibrium, the game collapses to pure price competition in the spirit of Varian (1980). Therefore, firms earn the same profit: $\frac{(1-\mu_L)r}{2}$ and price according to the same distribution function, independently of the timing structure.

4.3.2 Policies to Improve Consumer Sophistication

Economic theory emphasises the need to protect consumers by creating an effective competitive process rather than artificial supply-side controls. It is therefore appropriate to analyse the comparative statics of the model by considering the effect of introducing two policy measures that: (i) Improve the quantity of informed buyers, interpreted as an equal increase in μ_L , μ and μ_H , (ii) Reduce the marginal effect of obfuscation, interpreted as reducing the differences: $\mu_L - \mu$, $\mu - \mu_H$. In this section we also relax the assumption of a constant marginal effect of obfuscation and focus on the most interesting range: $\delta < \delta^c$.¹⁹

¹⁹It is most instructive to focus on the implications for firm obfuscation because consumer welfare is inherently sensitive to the specification of the utility function.

Proposition 4. The implications of policy measures are as follows:

1. Policies that increase the proportion of informed consumers lead to higher obfuscation by firms (λ_L decreases).
2. Policies that reduce the marginal effect of obfuscation may lead to higher or lower obfuscation by firms:
 - (i) As μ increases, λ_L increases
 - (ii) As μ_H increases, λ_L decreases.
 - (iii) As μ_L increases, λ_L decreases.

Policies that increase the amount of informed buyers lead to higher obfuscation by firms. The intuition is that firms seek to restore the original level of consumer sophistication to soften price competition. Therefore policies that affect the proportion of informed buyers may lead to more complex market interactions for consumers.

Gu and Wenzel (2014) also show that firms increase their obfuscation in response to policies that reduce the effectiveness of obfuscation but for a different reason. In their model, firms select obfuscation before setting prices and one firm is more prominent. The prominent seller receives a larger share of uninformed buyers independently of its price and obfuscation. They find that the prominent firm selects maximum obfuscation whilst the less prominent firm may not. Subsequently, policies to reduce the effectiveness of obfuscation may induce the less prominent firm to increase their obfuscation. In our model, firms are symmetric but maximum obfuscation ($\lambda_L = 0$) is sub-optimal because a more competitively priced firm benefits from a larger fraction of buyers being informed. A policy that reduces the effect of obfuscation on consumer confusion, however, incentivises both firms to increase their obfuscation to counteract the increase in consumer sophistication.

Armstrong et al. (2009) also show that supply-side policies may not benefit consumers but their explanation hinges on the incentives for consumers to engage with the market. They study maximum price limits and show that maximum prices prevent firms from charging the highest prices, which directly benefits consumers. However, the reduction in price dispersion reduces consumers' incentives to become informed. This softens competition amongst firms and can leave consumers paying higher prices.

Limiting the effectiveness of obfuscation on consumer sophistication creates mixed incentives for firms depending on the precise incidence on model parameters. We reconcile this mismatch of incentives in the following way. Increasing μ reduces the marginal effect of obfuscation because the difference, $\mu_L - \mu$, declines. Therefore, the marginal reduction in price competition from

obfuscating decreases, which makes obfuscation less attractive. However, policies that increase μ_H and μ_L incentivise firms to increase their obfuscation. The reason is that when μ_L and μ_H are higher, price competition will be more fierce when firms select low complexity. Therefore, obfuscation becomes more attractive to soften price competition.

These results are salient for competition authorities because improving consumer sophistication is often championed to make markets operate better for consumers. However, when consumers are simplicity biased, demand-side policies may lead to greater obfuscation by firms. We may also use these results to draw inter-industry comparisons. The model predicts that firms will obfuscate more in industries where consumers are more sophisticated.

4.3.3 An Illustrative Example

Consider a market where 75% of consumers correctly identify the lowest priced seller when firms do not obfuscate (μ_L).²⁰ This figure reduces to 50% when one firm obfuscates (μ) and 25% when both firms obfuscate (μ_H). For this constellation of parameters each firm assigns at least 50% probability to low complexity, depending on the degree of simplicity bias, as illustrated in Figure 4.2. For $\delta \in [0, 0.5)$ the weight assigned to low complexity increases linearly until low complexity becomes a pure strategy solution for $\delta \geq 0.5$.




Figure 4.2: Probability Assigned to Low Complexity

²⁰The idea that some consumers fail to select the cheapest seller when the range of options is relatively simple is consistent with empirical evidence (e.g. Wilson & Waddams-Price, 2010).

Suppose that 25% of uninformed consumers buy from the least complex seller ($\delta = 0.25$).²¹ The associated support defining prices are given by inserting the respective values into (4.6) and (4.7) to yield $p^l = 0.27$ and $p^m \approx 0.79$. $Supp F^L(p') = [0.27, 0.79]$ and $Supp F^H(p) = [0.79, 1]$.

Insert these values into (4.11) to obtain the probability that firms' assign to low complexity: $\lambda_L = 0.75$. That is, firms randomise between low and high complexity according to the ratio: 3:1. Firm profit is given by inserting the parameters into equation (4.14) to yield: $\pi_i = 0.23$. Industry profit is approximately 0.46, which generates a consumer surplus of 0.54. The probability distributions governing prices are given by inserting the respective parameters into expressions (4.15) and (4.16), as illustrated in Figure 4.3. The illustration also confirms the adjacency of the supports, the absence of any point masses and that both distributions are increasing in p .




Figure 4.3: Price Distribution ($\delta = 0.25$)




Figure 4.4: Price Distribution ($\delta = 0.4$)

Each firm places zero probability on prices below 0.27, irrespective of the complexity level it chooses. When the firm selects low complexity, prices are drawn from the range $[0.27, 0.79]$. When the firm selects high complexity, prices are drawn from the range $[0.79, 1]$. The probability of a tie in prices is zero, ensuring that the less complex firm always charges a lower price than the more complex firm.

Figure 4.4 illustrates the case where the degree of simplicity bias is larger and 40% of uninformed consumers purchase from the least complex seller ($\delta = 0.4$). The increase in the strength of the bias reduces p^l to 0.2 and increases p^m to 0.87, increasing price dispersion. The probability

²¹This satisfies the constraint for $\lambda_L < 1$. Any $\delta < 0.5$ would also satisfy this constraint. For $\delta \geq 0.5$, $\lambda_L = 1$ and firms are confined to low complexity. That is, the simplicity bias is effective in preventing firms from obfuscating.

that firms select low complexity also increases from 0.75 to 0.9 because the incentive to compete for simplicity biased consumers strengthens. This leads to a lower level of market complexity and a larger fraction of consumers are able to comprehend prices. Price competition intensifies, industry profit declines from 0.46 to 0.345 and consumer welfare increases by 0.115.

4.4 Endogenous Simplicity Biases

In the main model the magnitude of the simplicity bias is exogenous and reflects the type of consumer behaviour observed in retail markets and identified in experimental and empirical studies. A natural question is whether firms hold an incentive to stimulate consumers' biases towards simple products? To address this, we now consider a simple extension in which firms are able to influence the degree of bias using long-term marketing messages that encourage or dissuade consumers from selecting simple products, which endogenises the simplicity bias.

This is a specific kind of persuasive advertising that seeks to influence the way uninformed consumers arrive at a decision rather than persuading consumers to purchase from a particular seller. In particular, we subtly distinguish between marketing techniques that: (i) make a product more comprehensible for buyers, which would attract simplicity biased consumers and increase the proportion of informed buyers, (ii) encourage consumers to select the most simple option, which increases the degree of bias.

Piccione and Spiegler (2012) analyse a model in which firms select both a price and a marketing component. Their marketing element is interpreted as a price frame, which influences consumers' ability to compare prices. In our model firms choose a price, its complexity (framing in the terminology of Piccione and Spiegler (2012)) and a separate marketing message that influences consumers' choice procedures. Therefore, in addition to marketing influencing consumer choice, our model also allows for marketing that influences consumers' choice procedure.

4.4.1 Strategic Advertising

Consider a preliminary stage to the game in which each firm $i \in N$ makes a binary choice regarding their advertising message. Specifically, firms costlessly select between encouraging (E) or dissuading (D) consumers from selecting the simplest option in the market. The combination of firms' choices subsequently maps to a degree of simplicity bias $\delta_{A_i, A_j} \in (0, 1)$ where:

$$\delta_{E,E} > \delta_{E,D} = \delta_{D,E} > \delta_{D,D}$$

The decision to encourage or dissuade consumers from selecting the most simple option deliberately precedes the main game in which firms select a price and complexity level for their product,

for two reasons. Firstly, engineering an advertising campaign requires a longer lead time than adjusting the price or obfuscation of a product. Secondly, efforts to influence consumer decision-making require time to take effect. The following timeline defines the structure of the three-stage augmented game.

Period 1:

- Each firm $i \in N = \{1, 2\}$ selects (without cost) an advertising strategy $A_i \in \{E, D\}$, which maps to a degree of simplicity bias $\delta_{A_i, A_j} \in (0, 1)$.

Period 2:

- Each firm $i \in N = \{1, 2\}$ selects (without cost) a price p_i and a complexity level $k_i \in \{L, H\}$ for their product.
- The proportion of informed consumers is determined as a declining function of market complexity such that: $0 < \mu_H < \mu < \mu_L < 1$, where $\{L, L\} = \mu_L$, $\{L, H\} = \{H, L\} = \mu$, $\{H, H\} = \mu_H$.

Period 3:

- Consumers purchase according to their type:
 - $\mu_{k_i, k_j} \in (0, 1)$ are Informed and buy from firm with the lowest price.
 - $\delta_{A_i, A_j}(1 - \mu_{k_i, k_j})$ are Simplicity Biased and buy from the firm with the least complex promotion.
 - $(1 - \delta_{A_i, A_j})(1 - \mu_{k_i, k_j})$ are Random Purchasers and buy from a randomly selected firm.
- Firms receive their respective profit.

Assuming that advertising messages are binary is not restrictive and the results are preserved if advertising is a scalar variable. The extended game is solved using backward induction. First, consider the second period in which firms select both a price and a complexity level to determine the optimal firm behaviour for each contingency of the first period. Second, calculate the optimal advertising strategy of each firm in stage 1, conditional on the expected payoffs.

Proposition 5. There exists a unique pure strategy symmetric equilibrium in the first stage of the extended game in which all firms dissuade consumers from selecting the most simple option, before competing as in the main model.

It is initially surprising that firms do not seek to exacerbate consumers' behavioural biases. The intuition follows that in the absence of a simplicity bias, consumers who cannot observe prices are indifferent regarding the two sellers and purchase randomly. This removes the incentive for firms to compete in complexity and allows firms to obfuscate more. The increase in obfuscation maps to a corresponding increase in the proportion of uninformed consumers, which softens price competition and allows firms to earn higher profits.

This result is especially interesting because the main model predicts that firms should engineer promotional strategies that appeal to consumers who prefer a simple option. However, the extended model indicates that firms should avoid any promotional activity that would stimulate consumers' preferences for simple products. Therefore, an optimal promotional strategy should offer a simple product, to attract simplicity biased consumers, whilst encouraging consumers not to select the most simple option. This intuition can explain why firms with simple products do not directly encourage consumers to compare on the basis of complexity.

4.4.2 False Advertising

Firms may also be able to reduce consumers' perceptions of the complexity of their product or price scheme using marketing, as described in Chapter 3. Therefore, sellers may seek to prevent consumers from identifying the true 'simplest deal.' In the main model this is not possible because consumers correctly distinguish between firms that select high and low complexity.²² However, in practice firms may embed their product in a more elaborate marketing campaign that emphasises simplicity to prevent this distinction.

For example, *O₂* offer a 'Simplicity' tariff for mobile phones, which separates handset costs from monthly usage costs. They claim that this simplifies the product and allows consumers to purchase 'sim-only' plans. However, by unbundling their prices, consumers may encounter greater difficulties when interpreting and comparing tariffs. Other examples include complex financial products that are accompanied by a simple overview, which simultaneously limit comparability and appeal to consumers' biases. In our framework, this is equivalent to a reduction in the degree of simplicity bias: Fewer uninformed consumers buy directly from the firm with the lowest complexity.²³

This differs from the conventional interpretation of false advertising in which sellers advertise a characteristic that is factually incorrect (Rhodes and Wilson, 2015). Moreover, this line of thought suggests that the scope of false advertising should be widened and regulators should be weary

²²The weaker condition: That a majority of consumers identify which products are most complex, does not qualitatively affect our results.

²³A more extreme consideration might consider a negative simplicity bias if firms can direct consumers towards more complex products. However, this seems less plausible.

that firms may seek to simultaneously obfuscate their prices and products but appear transparent and simple to consumers, which to the best of our knowledge has not been identified before.

4.5 Discussion

This article enriches our understanding of competition when consumers exhibit a bias towards simple products. In particular, our results provide a plausible explanation for the observed progression towards simple products and promotion styles by retailers. In this section we discuss the main findings with reference to the related study in Chapter 3, draw comparisons with the wider literature, highlight the implications for policy and review the sharp predictions that lend particularly well to experimental and empirical study.

One of our main results is that the introduction of simplicity biased consumers counteracts the incentive for a firm to obfuscate. This leads to more transparent markets for consumers and enables more individuals to accurately compare prices. As a consequence, firms compete more fiercely for the demand of the larger population of informed consumers, which leads to lower average prices and lower profits.

Simple products are also least expensive in equilibrium because a firm that charges a high price benefits the most from confusing consumers. Therefore, consumers who choose the simple option because they cannot identify which firm is cheaper, inadvertently purchase from the cheapest seller. This confers an element of rationality to the seemingly heuristic-based choice procedure. The implications for firm profit are also interesting and differ significantly from the model in Chapter 3. In this study, simplicity biased consumers benefit directly from their bias in the form of lower prices (Direct effect) and incentivise firms to reduce their obfuscation, which stimulates price competition (Externality effect). Therefore, the direct and externality effects lead to a multiplied reduction in firms' profit.

In the sequential framework in Chapter 3, firms attach a price premium to simple products (Homburg et al, 2014). The intuition follows that the firm with the least complex promotion is guaranteed a larger share of the uninformed consumers than his rival. Therefore, any reduction in price to compete for informed consumers in the second stage, is more costly for the firm with the simple promotion. This generates a negative correlation between price and obfuscation.

In this sense, the direct and externality effects are competing forces in the sequential setting but complementary effects on firm profit in this simultaneous environment, generating different policy implications. Therefore, these studies provide a plausible explanation for heterogeneous correlations between price or product complexity and prices. Interestingly, when the bias is sufficient to generate a low complexity equilibrium in the sequential model, the pricing behaviour of

firms is also independent of the timing structure. Both firms price randomly from the same distribution, $F^L(p)$, and earn the same payoff that depends only on the 'natural' level of consumer sophistication (μ_L).

The simultaneous and sequential models are also united by a prediction of stochastic pricing and stochastic obfuscation behaviour, providing the degree of bias is not sufficiently small (large) to obtain a high (low) complexity equilibrium. However, the frameworks diverge in the characteristic properties of the price distributions. In the simultaneous game (Chapter 4), firms draw prices according to non-atomic and continuous probability functions. This prevents the rival from undercutting by a marginal deviation to steal the demand of informed buyers. In this sense, firms continually engage in sale pricing. However, in the sequential framework (Chapter 3), the price distribution of the least complex seller contains a mass point at the headline price (r) whenever firms differ in their obfuscation. This implies that less complex firms will vary between the recommended retail price and sale prices but the more complex firm always chooses sale prices.

The supports of the price distributions are also mutually exclusive in the simultaneous framework as firms always combine higher prices with greater obfuscation. However, in the sequential framework, the support of the price distribution adopted by the more complex firm is nested in the support of the price distribution of the less complex firm. To test these models we must therefore examine average prices, in addition to individual prices. This further supports the idea that uninformed consumers may infer that more complex prices are always more expensive in markets where complexity is quickly adjustable but fail to identify that simple products are more expensive on average when complexity is a long-term parameter. These models also provide a credible explanation for the observed heterogeneity in pricing strategies across firms and across retail industries. Furthermore, the introduction of simplicity biases in consumer decision-making provides a novel explanation for sellers to frequently change the complexity of their promotions, in addition to their price.

The model we study in this chapter complements and extends the underlying duopoly framework of Chioveanu and Zhou (2013) with a revised focus on the impact of simplicity biases and the introduction of endogenous biases due to marketing. In particular, the duopoly model they analyse acts as a benchmark for the present model when the degree of simplicity bias is zero and exogenous. It is therefore not surprising that some of the equilibrium properties are common to both models, including the positive correlation between prices and obfuscation. The most significant divergence between the outcomes of the models lies in the obfuscation strategy for firms. More specifically, firms will always obfuscate less in the present study and perhaps not at all, because simplicity biased consumers counteract the incentive for firms to obfuscate.

The intuition for the mixed strategy in obfuscation is also different. In their model, firms randomise between bundles of high prices, high complexity and low prices, low complexity because

a more competitively priced firm benefits from a larger quantity of informed clients. In this sense, obfuscation exists only as a tool to manipulate the composition of the consumer population in terms of informed and uninformed buyers. However, in our model, firms also vary their obfuscation to compete for simplicity biased consumers. By exploring the incidence of simplicity biases further and endowing firms with the ability to influence consumers' biases using marketing, we are also able to explain why firms actively advertise the simplicity of their product but do not directly encourage consumers to purchase on the basis of simplicity.

This explanation for firms to frequently change the complexity of their product is also distinct from alternative explanations based on consumer learning. Carlin and Manso (2011) explore a continuous time model in which the proportion of informed consumers increases over time and firms can reset the learning process by changing the menu of their offer. They derive the optimal period of time that should elapse before a firm refreshes its menu to the point where the cost of revising the menu is exactly offset by the difference between the rent achieved at the beginning of the period and the rent available with the current lower proportion of informed consumers.

Interestingly, Carlin and Manso (2011) show that the relationship between the refresh rate and the speed of consumer learning is non-monotonic. Therefore, education initiatives that accelerate consumer learning may lead firms to refresh their menu more frequently, depending on the initial rate of consumer learning. Consumer learning is not directly modelled in our framework but one can interpret the simplicity bias as a form of learning if confused consumers begin to infer that simple products are cheaper. In our incentives structure, however, firms change not only their menu, which resets the learning process in Carlin and Manso (2011) but also the complexity of their menu as they compete for simplicity biased consumers.

The results also support the recent literature on common standards in retail markets. In that literature, obfuscation involves firms departing from industry-wide practices to make their product less comparable. Gaudeul and Sugden (2012) describe how firms are prevented from deviating to individuated promotion formats if consumers automatically exclude such sellers when comparing prices. In this sense, their model captures an extreme simplicity bias in which consumers exclude complex options before selecting from the remaining sellers. In our model, only some consumers exhibit a simplicity bias, which explains why firms in certain industries adhere to common standards whilst firms in other industries do not.

Moreover, in our model, individuated promotions are only one method of many that are available to a firm when devising an obfuscation strategy. For example, consumers may find comparisons more difficult if both firms adopt high complexity than low complexity, even though they follow the same price frame structure. In this sense, our focus relates to a broader class of strategies and binds together the recent literature on obfuscation, in which sellers gratuitously

add complexity to their products, with the emerging literature on common standards, in which sellers voluntarily conform to industry standards.

4.5.1 Policy Implications

Three main policy implications arise from this study. First, policies to remove biases in consumer decision-making may lead to lower consumer welfare and a weaker competitive process. Policies to discourage consumers from selecting the simplest product allow firms to obfuscate more and charge higher prices, which reduces consumer welfare if they return to purchasing randomly. Conversely, policies that encourage uninformed consumers to confine their choice set to less complex options improve consumer welfare. The model also demonstrates that the need for policies to prevent obfuscation may be misplaced as demand-side forces may endogenously deter such practices. This reinforces the idea that firms will voluntarily comply with common standards to maximise their potential clients (Gaudeul and Sugden, 2012).

Second, policies to enhance consumer sophistication, such as price comparison websites or decision-making aids, may induce adverse incentives for firms and fail to benefit consumers. Policies to improve the natural level of consumer sophistication (μ_L) cause firms to increase their obfuscation. Critically, this result does not hinge on ex ante asymmetric firms as described by Gu and Wenzel (2014). Therefore, demand-side remedies to market complexity must consider supply-side incentives. Action specific policies that specify firms' promotion structures may also be less effective in enhancing market competition than anticipated, if the underlying incentive to obfuscate remains and firms are able to create consumer confusion using alternative dimensions of obfuscation.

Third, regulators should assess the extent to which complex products being marketed as simple options constitutes false advertising. By introducing complexity, firms stimulate consumer confusion. However, if firms shroud their complexity behind a framework of marketing, simplicity biased consumers may actually purchase the most complex and most expensive products.

4.5.2 Hypotheses for Experimental & Empirical Research

The analysis presented in this article yields several striking predictions that are experimentally and empirically testable. Firstly, the model predicts that firms will frequently revise the complexity of their promotions in addition to the familiar randomisation in prices. This is an interesting avenue for empirical research as we often observe firms changing their promotion formats, product bundles and advertising campaigns more frequently than one might explain using consumer learning arguments.²⁴ Furthermore, empirical research should seek to identify whether

²⁴For example, firms often change their promotions when the product is not a repeated purchase for customers.

firms change their complexity or prices in a systematic method, which would expose the firm to a rival.

Secondly, obfuscation should be lower and prices more dispersed in markets where the degree of simplicity bias is more pronounced. This may explain the observation by Miravete (2013) that firms operating in the US cellular telecoms industry did not respond to the introduction of competition by increasing their obfuscation. In this sector, consumers may exhibit strong preferences towards simple products, which prevent firms from increasing their obfuscation.

Thirdly, in conjunction with the model presented in Chapter 3, this thesis highlights a distinction between short and long term obfuscation techniques. When complexity is a short-term parameter, the most complex firm should charge the highest price. When obfuscation is a longer term commitment that firms decide before competing in prices, the least complex firm charges the highest price. Empirical analysis may seek to categorise obfuscation methods according to their flexibility, to test these predictions directly. Fourthly, sellers may attempt to combine complex products and promotions with features that attract simplicity seeking consumers. Finally, obfuscation should be highest in markets where consumers would otherwise be most sophisticated.

4.5.3 Limitations & Extensions

The generality of this framework offers a simple mechanism that is sufficiently detailed to understand the impact of simplicity biases on market competition. However, there are several opportunities for development. In particular, obfuscation is binary as firms choose between low or high complexity. In practice, firms may be able to adopt a wider range of complexity levels in promoting their product. The reason that a continuous complexity scale is challenging is that the payoff to a firm now depends on both the obfuscation choice of the rival and the difference between the firms' obfuscation levels. Furthermore, it is clear that consumers will distinguish between simple and complex promotions, but this distinction may become blurred if obfuscation is a continuous variable and firms choose similar obfuscation levels.

The model also only considers two firms to highlight the role of simplicity biased consumers clearly. The main difficulty in an oligopoly model is the consumer decision procedure and the mechanism by which consumers become impaired by additional firms obfuscating. These topics lend themselves particularly well to further examination in an experimental setting, before being incorporated into a theoretical model of competition.

The degree of bias is also exogenous in the main model and a function of firms' marketing messages in the extension. The origins and factors that influence the degree of bias in consumer decision-making, however, remain relatively unexplored. In addition, little is known regarding

consumer demographics and systematic characteristics of individuals that are most likely to exhibit a simplicity bias. These are salient if authorities consider the distributional implications of simplicity biases.

The provenance of simplicity biases is also likely to vary across industries and interact with other behavioural tendencies of consumers, including *status quo* biases and inattention. Exploring these topics is of significant interest to ensure that theoretical frameworks are attuned to consumer behaviour, and regulatory policies are considerate to consumers' biases and non-standard approaches to decision-making. Future research might also incorporate simplicity biased consumers into models of multi-product retailing and consider the implications for information intermediaries.

4.6 Conclusion

In this article we have demonstrated that the the incentive for firms to obfuscate is significantly weaker when consumers are biased towards simple products. As a consequence markets become more transparent and more consumers are able to understand and compare prices. This leads to more fierce price competition, greater price dispersion and higher consumer welfare. The simplicity bias is also rational for uninformed consumers because in equilibrium more complex products are more expensive. Policies to reduce demand-side behavioural biases may not be welfare enhancing because such policies allow firms to obfuscate more and charge higher prices. Policies to improve consumer sophistication and resilience to obfuscation may also incentivise firms to obfuscate more, without requiring heterogeneities across firms.

In the extended advertising game, we distinguish between obfuscation and advertising. We show that firms may adopt a simple promotion scheme to attract biased consumers but subtly firms should not encourage consumers to select the simplest option, which would lead to greater price competition and erode industry profits. In addition, firms may attempt to circumvent the competitive forces of simplicity biases in consumers' choice procedures by falsely advertising complex products as simple options. Cumulatively, this article provides a detailed treatment of competition when consumers prefer less complex products, whilst motivating significant scope for further research.

4.7 Appendix A: Main Proofs

Proof of Lemma 1: The reservation price must form the upper bound of the price distribution associated with at least one level of complexity. For any maximum price \bar{p} below the reservation price r , the firm could increase its price to r with no reduction in demand and an increase in firm profit, contradicting \bar{p} as the equilibrium upper bound.

It remains to show that a firm will combine a lower price with low complexity for $\delta < \delta^c$. Let the expected payoff to firm i , given price-complexity profile (p_i, k_i) , be written $\pi_i(p_i, k_i)$. We must now show $\pi_i(H, 1) \geq \pi_i(L, 1)$.

$$\pi_i(H, 1) = \frac{1}{2} \cdot [\lambda_L \cdot (1 - \delta)(1 - \mu) + \lambda_H \cdot (1 - \mu_H)]$$

$$\pi_i(L, 1) = \frac{1}{2} \cdot [\lambda_L \cdot (1 - \mu_L) + \lambda_H \cdot (1 + \delta)(1 - \mu)]$$

This translates into the following condition:

$$\lambda_L \cdot (1 - \delta)(1 - \mu) + \lambda_H \cdot (1 - \mu_H) \geq \lambda_L \cdot (1 - \mu_L) + \lambda_H \cdot (1 + \delta)(1 - \mu)$$

$$\lambda_L \cdot (\mu_L - \mu - \delta + \delta\mu) + \lambda_H \cdot (\delta\mu - \delta + \mu - \mu_H) \geq 0$$

The condition is satisfied for any $\delta \in [0, \delta^c]$ and $\lambda_L \in [0, 1]$. Note that the assumption of a constant marginal effect of obfuscation implies that $\delta < \frac{\mu_L - \mu}{1 - \mu}$ is the same condition as $\delta < \frac{\mu - \mu_H}{1 - \mu}$. ■

Proof of Proposition 1: (Short Version).

Step 1: Demonstrate $\pi_i(H, p) \geq \pi_i(L, p)$ for all p in the support of $F^H(p)$. This requires:

$$\lambda_L \cdot (1 - \delta)(1 - \mu) + \lambda_H \cdot (1 + \mu_H [1 - 2F^H(p)]) \geq$$

$$\lambda_L \cdot (1 - \mu_L) + \lambda_H \cdot (1 + \delta + \mu [1 - \delta - 2F^H(p)])$$

Simplify and insert the definitions of λ_L and λ_H given in (4.11) to yield:

$$\frac{\mu - \mu_H + \delta - \delta\mu}{\mu_L - \mu_H} \cdot (\mu_L - \mu + \delta\mu - \delta) \geq \frac{\mu_L - \mu + \delta\mu - \delta}{\mu_L - \mu_H} \cdot (\delta - \delta\mu + (\mu - \mu_H) [1 - 2F^H(p)])$$

By removing the common factors and simplifying terms, this condition collapses to:

$$1 \geq [1 - 2F^H(p)] \implies F^H(p) \geq 0$$

By definition of a cumulative probability distribution function this condition is satisfied.

Step 2: Demonstrate $\pi_i(L, p') \geq \pi_i(H, p')$ for all p' in the support of $F^L(p')$. This requires:

$$\begin{aligned} \lambda_L \cdot (1 + \mu_L [1 - 2F^L(p')]) + \lambda_H \cdot (1 + \delta + \mu - \delta\mu) \geq \\ \lambda_L \cdot (1 - \delta + \mu [1 + \delta - 2F^L(p')]) + \lambda_H \cdot (1 + \mu_H) \end{aligned}$$

Simplify and insert the definitions of λ_L and λ_H :

$$\frac{\mu - \mu_H + \delta - \delta\mu}{\mu_L - \mu_H} \cdot (\delta - \delta\mu + (\mu_L - \mu) [1 - 2F^L(p')]) \geq \frac{\mu_L - \mu + \delta\mu - \delta}{\mu_L - \mu_H} \cdot (\mu_H - \mu + \delta\mu - \delta)$$

Once simplified this condition collapses to:

$$1 \geq F^L(p')$$

By definition of a cumulative probability distribution function the condition is satisfied. This concludes the proof that $Supp F^H(p) = [p^m, 1]$ and $Supp F^L(p') = [p^l, p^m]$. ■

Proof of Proposition 2: The proof proceeds in steps. First, we show that whenever the support of the price distribution associated with low complexity lies directly below the price distribution associated with high complexity, low complexity is a pure strategy equilibrium. Second, we show if the lower price distribution is instead associated with high complexity, there exists no profitable deviation from low complexity. Therefore, for any price firms select low complexity as a unique pure strategy equilibrium.

Step 1: Assume, as before, that the support of $F^L(p')$ lies directly below the support of $F^H(p)$. This implies a firm that selects low complexity always charges a price below the high complexity firm because a tie occurs with probability zero. We must now show that:

1. $\pi_i(L, p') \geq \pi_i(H, p')$

$$2. \pi_i(L, p) \geq \pi_i(H, p)$$

That is, the payoff associated with low complexity is not smaller than the payoff associated with high complexity, for any price in the supports of the respective distributions. The first condition is identical to the the main model and continues to hold for any $\delta \in [0, 1]$:

$$\begin{aligned} \lambda_L \cdot [1 + \mu_L(1 - 2F^L(p'))] + \lambda_H \cdot (1 + \delta + \mu - \delta\mu) \geq \\ \lambda_L \cdot [1 - \delta + \delta\mu + \mu(1 - 2F^L(p'))] + \lambda_H \cdot [1 + \mu_H] \end{aligned}$$

$$\lambda_L \cdot [\delta - \delta\mu + (\mu_L - \mu)(1 - 2F^L(p'))] + \lambda_H \cdot [\delta + \mu - \delta\mu - \mu_H] \geq 0$$

The condition is satisfied for all $\lambda_L \in [0, 1]$ because $\delta \geq \delta^c$. The second condition is the direct opposite of the condition in the main model and states that a firm will prefer to combine a price from the high complexity distribution with low complexity:

$$\begin{aligned} \lambda_L \cdot (1 - \mu_L) + \lambda_H \cdot (1 + \delta - \delta\mu + \mu [1 - 2F^H(p)]) \geq \\ \lambda_L \cdot (1 - \delta)(1 - \mu) + \lambda_H \cdot (1 + \mu_H [1 - 2F^H(p)]) \end{aligned}$$

$$\lambda_L \cdot [\delta - \delta\mu - \mu_L + \mu] + \lambda_H \cdot [\delta - \delta\mu + (\mu - \mu_H)[1 - 2F^H(p)]] \geq 0$$

The condition is satisfied for any $\lambda_L \in [0, 1]$ $\because \delta \geq \frac{\mu_L - \mu}{1 - \mu} = \frac{\mu - \mu_H}{1 - \mu}$. Cumulatively these results demonstrate that, conditional on a low complexity firm charging a price below a high complexity firm, there exists no profitable deviation away from the strategy $(L, F^L(p'))$.

Step 2: It remains to show that if a firm selects high complexity and charges a lower price than a firm that selects low complexity, profit is also below the level associated with low complexity. Formally, suppose that the support of $F^H(p)$ lies below $F^L(p')$. A firm that selects high complexity is now more expensive than a firm that selects low complexity with probability 1.

We first show that the profile $(L, 1)$ continues to be an equilibrium when $Supp F^H(p) \leq Supp F^L(p')$. It follows from standard arguments that the equilibrium price distribution must be bounded from above at the reservation threshold of the consumer population, which is 1. Therefore, we require $\pi_i^R(L, 1) \geq \pi_i^R(H, 1)$. Superscript R indicates reversed supports. This corresponds to:

$$\lambda_L \cdot [1 - \mu_L] + \lambda_H \cdot [(1 + \delta)(1 - \mu)] \geq \lambda_L \cdot [(1 - \delta)(1 - \mu)] + \lambda_H \cdot [1 - \mu_H]$$

$$\lambda_L \cdot [\delta - \delta\mu + \mu - \mu_L] + \lambda_H \cdot [\delta - \delta\mu + \mu_H - \mu] \geq 0$$

This condition is always satisfied for $\delta \geq \delta^c$. We are now able to show that when $Supp F^L(p') \geq Supp F^H(p)$, there exists no profitable deviation in complexity from Low. This requires:

1. $\pi_i^R(L, p) \geq \pi_i^R(H, p)$
2. $\pi_i^R(L, p') \geq \pi_i^R(H, p')$.

The first condition requires:

$$\begin{aligned} \lambda_L \cdot [1 + \mu_L] + \lambda_H \cdot [1 + \delta - \delta\mu + \mu[1 - 2F^H(p)]] &\geq \\ \lambda_L \cdot [1 + \mu + \delta\mu - \delta] + \lambda_H \cdot [1 + \mu_H(1 - 2F^H(p))] & \end{aligned}$$

$$\lambda_L \cdot [\mu_L + \delta - \mu - \delta\mu] + \lambda_H \cdot [\delta - \delta\mu + (\mu - \mu_H)(1 - 2F^H(p))] \geq 0$$

The second condition requires:

$$\begin{aligned} \lambda_L \cdot [1 + \mu_L(1 - 2F^L(p'))] + \lambda_H \cdot [(1 + \delta)(1 - \mu)] &\geq \\ \lambda_L \cdot [1 - \delta + \delta\mu + \mu(1 - 2F^L(p'))] + \lambda_H \cdot [1 - \mu_H] & \end{aligned}$$

$$\lambda_L \cdot [\delta - \delta\mu + (\mu_L - \mu)(1 - 2F^L(p'))] + \lambda_H \cdot [\delta - \delta\mu + \mu_H - \mu] \geq 0$$

Both conditions are satisfied for any $\delta \geq \delta^c$. Therefore, irrespective of the price, firm profit is always higher when low complexity is chosen for $\delta > \frac{\mu_L - \mu}{1 - \mu}$.

To derive the price distribution recognise that the combination $(L, 1)$ is now an element of the mixed strategy equilibrium with an associated payoff: $\frac{1 - \mu_L}{2}$. For any price p' in the support of the symmetric price distribution, the expected payoff is: $\left[\frac{1 - \mu_L}{2} + [1 - F^L(p')] \mu_L \right] p'$. This yields the following indifference condition:

$$\frac{1 - \mu_L}{2} = \left[\frac{1 - \mu_L}{2} + [1 - F^L(p')] \mu_L \right] \cdot p'$$

which solves to:

$$F^L(p') = 1 - \frac{1 - \mu_L}{2\mu_L} \left(\frac{1 - p'}{p'} \right)$$

with support: $\left[\frac{1-\mu_L}{1+\mu_L}, 1\right]$. ■

Proof of Proposition 3:

(1i) The proof is given immediately by differentiating the probability function w.r.t δ :

$$\lambda_L = \frac{\mu - \mu_H + \delta - \delta\mu}{\mu_L - \mu_H} \implies \frac{\partial \lambda_L}{\partial \delta} = \frac{1 - \mu}{\mu_L - \mu_H} > 0$$

An increase in the degree of simplicity bias causes firms to assign larger probability to low complexity. ■

(1ii) Differentiate the expression for firm profit with respect to δ . This yields:

$$\pi_i = \frac{1}{2} \cdot \left[1 - \mu_H - \frac{(\mu - \mu_H + \delta - \delta\mu)^2}{\mu_L - \mu_H} \right]$$

$$\frac{\partial \pi_i}{\partial \delta} = - \frac{(\mu - \mu_H + \delta - \delta\mu)(1 - \mu)}{\mu_L - \mu_H} < 0$$

Consumer surplus (CS) is the reciprocal of industry profit, which is twice the value of firm profit in this symmetric game.

$$CS = 1 - \sum_{i \in N} \pi_i = \mu_H + \frac{(\mu - \mu_H + \delta - \delta\mu)^2}{\mu_L - \mu_H}$$

$$\frac{\partial CS}{\partial \delta} = \frac{2(\mu - \mu_H + \delta - \delta\mu)(1 - \mu)}{\mu_L - \mu_H} > 0$$
■

(1iii) We now show that p^l is decreasing in δ .

Step 1: The definition of p^l is given by:

$$p^l = \frac{\lambda_L \cdot (1 - \delta)(1 - \mu) + \lambda_H \cdot (1 - \mu_H)}{\lambda_L \cdot (1 + \mu_L) + \lambda_H \cdot (1 + \delta + \mu - \delta\mu)}$$

The numerator and denominator must both be positive in the relevant range. Define the numerator and denominator as $f(\delta)$ and $g(\delta)$ respectively. To demonstrate that $\frac{\partial p^l}{\partial \delta} < 0$ where $f(\delta), g(\delta)$ are both positive, it is sufficient to show that $f'(\delta) < 0$, $g'(\delta) > 0$ such that $f'g - fg' < 0$.

Step 2: Insert the expressions for λ_L and λ_H and remove the common denominator $(\mu_L - \mu_H)$ from all elements:

$$p^l = \frac{(\mu - \mu_H + \delta - \delta\mu) \cdot (1 - \delta)(1 - \mu) + (\mu_L - \mu + \delta\mu - \delta) \cdot (1 - \mu_H)}{(\mu - \mu_H + \delta - \delta\mu) \cdot (1 + \mu_L) + (\mu_L - \mu + \delta\mu - \delta) \cdot (1 + \delta + \mu - \delta\mu)}$$

Simplify the numerator :

$$p^l = \frac{(\delta - \delta\mu)(2\mu_H - 2\mu + \delta\mu - \delta) + (\mu - \mu_H)(1 - \mu) + (\mu_L - \mu)(1 - \mu_H)}{(\mu - \mu_H + \delta - \delta\mu) \cdot (1 + \mu_L) + (\mu_L - \mu + \delta\mu - \delta) \cdot (1 + \delta + \mu - \delta\mu)}$$

Step 3: We are now able to demonstrate that $f' < 0$ and $g' > 0$ by differentiating the numerator and denominator individually, which coupled with $f > 0$ and $g > 0$ implies $f'g - fg' < 0$.

$$f' = \frac{\partial f}{\partial \delta} = 2 \cdot (1 - \mu)(\mu_H - \mu + \delta\mu - \delta) < 0$$

The final term is always negative because $\mu > \mu_H$ and $\delta > \delta\mu$. This confirms that $f'g < 0$. Differentiating the denominator ($g(\delta)$) yields:

$$g' = \frac{\partial g}{\partial \delta} = 2(1 - \mu)(\mu_L - \mu + \delta\mu - \delta)$$

Using the restriction placed on the degree of simplicity bias: $\delta < \frac{\mu_L - \mu}{1 - \mu} \implies \mu_L - \mu + \delta\mu - \delta > 0$. This condition guarantees that $fg' \geq 0$. Therefore: $f'g - fg' < 0$. ■

(2) In the simultaneous game firms are confined to low complexity whenever:

$$\delta \geq \frac{\mu_L - \mu}{1 - \mu}$$

In the sequential case firms are disciplined to low complexity whenever:

$$\frac{1 - \mu_L}{1 - \mu} \geq \frac{1 - \delta^2 + \mu(1 + 2\delta + \delta^2)}{1 + \delta + \mu - \delta\mu}$$

Subtract 1 from each side of the expression above and simplify:

$$\frac{\delta^2 + \delta - 3\delta\mu - \delta^2\mu}{1 + \delta + \mu - \delta\mu} \geq \frac{\mu_L - \mu}{1 - \mu}$$

It is now possible to show that for any combination of $\mu \in (0, 1)$, $\delta \in (0, 1)$:

$$\delta \geq \frac{\delta^2 + \delta - 3\delta\mu - \delta^2\mu}{1 + \delta + \mu - \delta\mu} \implies 4\delta\mu \geq 0$$

Therefore, firms are confined to low complexity for a smaller degree of simplicity bias in the simultaneous game than in the sequential game. That is, whenever:

$$\frac{\delta^2 + \delta - 3\delta\mu - \delta^2\mu}{1 + \delta + \mu - \delta\mu} \geq \frac{\mu_L - \mu}{1 - \mu}$$

It must also be the case that:

$$\delta \geq \frac{\mu_L - \mu}{1 - \mu}$$

■

(3) In the sequential model firms earn a payoff equal to $\frac{1-\mu_H}{2}$, which is strictly larger than the simultaneous game payoff: $\frac{1}{2} \cdot \left[1 - \mu_H - \frac{(\mu - \mu_H)^2}{\mu_L - \mu_H}\right]$.

■

(4) The proof is given in the main text.

Proof of Proposition 4:

For part (1): $\lambda_L = \frac{\mu - \mu_H + \delta - \delta\mu}{\mu_L - \mu_H}$. A policy to increase consumer sophistication for every level of market complexity, that is: to increase μ_L, μ, μ_H by a fixed amount, causes λ_L to increase. To see this note that in the specification of μ_L , an increase in μ_L, μ, μ_H by a fixed quantity has no impact on the difference $\mu - \mu_H$ or $\mu_L - \mu_H$. Therefore, the only element that is affected is the RHS of the numerator: $\delta - \delta\mu$, which decreases.

For part (2): The result follows directly by differentiating λ_L w.r.t μ_L, μ, μ_H .

$$\frac{\partial \lambda_L}{\partial \mu_L} = -1 \cdot \frac{\mu - \mu_H + \delta - \delta\mu}{(\mu_L - \mu_H)^2} < 0$$

$$\frac{\partial \lambda_L}{\partial \mu} = \frac{(1 - \delta)(\mu_L - \mu_H)}{(\mu_L - \mu_H)^2} > 0$$

$$\frac{\partial \lambda_L}{\partial \mu_H} = \frac{-\mu_L + \mu + \delta - \delta\mu}{(\mu_L - \mu_H)^2} < 0$$

Note that in 4.7: $-\mu_L + \mu + \delta - \delta\mu < 1$ because we focus on the most interesting range where $\delta \leq \frac{\mu - \mu_H}{1 - \mu}$ and the decreasing (or non-increasing) marginal effect of obfuscation implies $\mu_L - \mu \geq \mu - \mu_H$ such that $\delta \leq \frac{\mu_L - \mu}{1 - \mu}$.

■

Proof of Proposition 5: The second stage mirrors the main model and the preceding equilibria are preserved. Therefore, the expected profit for each firm in the second stage is written:

$$\pi_i = \frac{1}{2} \left[1 - \mu_H - \frac{(\mu - \mu_H + \delta_{A_i, A_j}(1 - \mu))^2}{\mu_L - \mu_H} \right]$$

Firm profit is decreasing in the degree of simplicity bias. Therefore, all firms unilaterally choose to suppress the simplicity bias and encourage confused consumers to select randomly.

■

4.8 Appendix B: Supplementary Materials.

This section contains detailed derivations of the proofs in Appendix A and supplementary mathematical derivations for referees.

Proof of Equality between p^m Expressions: The expression for p^m may differ in exposition depending on which indifference condition is applied. However, they are equivalent as indicated below:

$$\frac{\lambda_L \cdot (1 - \delta)(1 - \mu) + \lambda_H \cdot (1 - \mu_H)}{\lambda_L \cdot (1 - \mu_L) + \lambda_H \cdot (1 + \mu + \delta - \delta\mu)} = \frac{\lambda_L \cdot (1 - \delta)(1 - \mu) + \lambda_H \cdot (1 - \mu_H)}{\lambda_L \cdot (1 - \delta)(1 - \mu) + \lambda_H \cdot (1 + \mu_H)}$$

$$\lambda_L \cdot (1 - \mu_L) + \lambda_H \cdot (1 + \mu + \delta - \delta\mu) = \lambda_L \cdot (1 - \delta)(1 - \mu) + \lambda_H \cdot (1 + \mu_H)$$

Rearrange and apply $\lambda_L = 1 - \lambda_H$ to yield:

$$\lambda_L(\mu_H - \mu_L) = \mu_H - \mu + \delta\mu - \delta$$

Insert the definition of λ_L from (4.11):

$$\frac{\mu - \mu_H + \delta - \delta\mu}{\mu_L - \mu_H} \cdot (\mu_H - \mu_L) = \mu_H - \mu + \delta\mu - \delta$$

$$\mu_H - \mu + \delta\mu - \delta = \mu_H - \mu + \delta\mu - \delta$$

■

Proof of Optimal Obfuscation in Proposition 1 (Detailed Version).

Step 1: Demonstrate $\pi_i(H, p) \geq \pi_i(L, p)$. This requires:

$$\lambda_L \cdot (1 - \delta)(1 - \mu) + \lambda_H \cdot (1 + \mu_H [1 - 2F^H(p)]) \geq$$

$$\lambda_L \cdot (1 - \mu_L) + \lambda_H \cdot (1 + \delta + \mu [1 - \delta - 2F^H(p)])$$

$$\lambda_L \cdot (1 - \delta - \mu + \delta\mu - 1 + \mu_L) \geq \lambda_H \cdot (1 + \delta - \delta\mu - 1 + (\mu - \mu_H) [1 - 2F^H(p)])$$

$$\lambda_L \cdot (\mu_L - \mu + \delta\mu - \delta) \geq \lambda_H \cdot (\delta - \delta\mu + (\mu - \mu_H) [1 - 2F^H(p)])$$

Insert the definitions of λ_L and λ_H :

$$\frac{\mu - \mu_H + \delta - \delta\mu}{\mu_L - \mu_H} \cdot (\mu_L - \mu + \delta\mu - \delta) \geq \frac{\mu_L - \mu + \delta\mu - \delta}{\mu_L - \mu_H} \cdot (\delta - \delta\mu + (\mu - \mu_H) [1 - 2F^H(p)])$$

$$(\mu - \mu_H + \delta - \delta\mu) \cdot (\mu_L - \mu + \delta\mu - \delta) \geq (\mu_L - \mu + \delta\mu - \delta) \cdot (\delta - \delta\mu + (\mu - \mu_H) [1 - 2F^H(p)])$$

Identify $(\mu_L - \mu + \delta\mu - \delta)$ as a (positive) common factor and remove. Positivity is implied $\delta < \delta^c$.

$$(\mu - \mu_H) + (\delta - \delta\mu) \geq (\delta - \delta\mu) + (\mu - \mu_H) [1 - 2F^H(p)]$$

$$1 \geq [1 - 2F^H(p)] \implies F^H(p) \geq 0$$

By definition of a cumulative probability distribution function $F^H(p) \geq 0$. This concludes the first element of the proof.

Step 2: Demonstrate $\pi_i(L, p') \geq \pi_i(H, p')$. This requires:

$$\begin{aligned} \lambda_L \cdot (1 + \mu_L [1 - 2F^L(p')]) + \lambda_H \cdot (1 + \delta + \mu - \delta\mu) &\geq \\ \lambda_L \cdot (1 - \delta + \mu [1 + \delta - 2F^L(p')]) + \lambda_H \cdot (1 + \mu_H) & \end{aligned}$$

$$\lambda_L \cdot (\delta - \delta\mu + (\mu_L - \mu) [1 - 2F^L(p')]) \geq \lambda_H \cdot (\mu_H - \mu + \delta\mu - \delta)$$

Insert the definitions of λ_L and λ_H .

$$\frac{\mu - \mu_H + \delta - \delta\mu}{\mu_L - \mu_H} \cdot (\delta - \delta\mu + (\mu_L - \mu) [1 - 2F^L(p')]) \geq \frac{\mu_L - \mu + \delta\mu - \delta}{\mu_L - \mu_H} \cdot (\mu_H - \mu + \delta\mu - \delta)$$

$$(\mu - \mu_H + \delta - \delta\mu) \cdot (\delta - \delta\mu + (\mu_L - \mu) [1 - 2F^L(p')]) \geq (\mu_L - \mu + \delta\mu - \delta) \cdot (\mu_H - \mu + \delta\mu - \delta)$$

Rewrite to identify common factor:

$$(\mu - \mu_H + \delta - \delta\mu) \cdot (\delta - \delta\mu + (\mu_L - \mu) [1 - 2F^L(p')]) \geq (-\mu_L + \mu - \delta\mu + \delta) \cdot (\mu - \mu_H + \delta - \delta\mu)$$

$$\delta - \delta\mu + (\mu_L - \mu) [1 - 2F^L(p')] \geq -\mu_L + \mu - \delta\mu + \delta$$

$$(\mu_L - \mu) [1 - 2F^L(p')] \geq -\mu_L + \mu$$

$$(\mu_L - \mu) [2 - 2F^L(p')] \geq 0 \implies 1 \geq F^L(p')$$

By definition $(\mu_L - \mu) > 0$. By definition of a cumulative probability distribution function the condition is satisfied. This concludes the proof that $Supp F^H(p) = [p^m, 1]$ and $Supp F^L(p') = [p^l, p^m]$. ■

Proof of Proposition 2: Supplementary Derivations

In the proof of Proposition 2, we show that when the supports are reversed, such that the support associated with high complexity lies below the support associated with low complexity, it remains optimal for a firm to select low complexity. In this supplementary component we provide the derivations of the payoffs used in the case where the supports are reversed. The payoffs in this case are denoted with superscript R to prevent confusion with the payoffs under the equilibrium supports of the price distributions.

$$\pi_i^R(H, p) = \left(\lambda_L \left[\mu + \frac{(1-\delta)(1-\mu)}{2} \right] + \lambda_H \left[\mu_H [1 - F^H(p)] + \frac{1-\mu_H}{2} \right] \right) p$$

When the firm selects high complexity under reversed supports: If the rival selects low complexity, the firm wins all of the informed consumers, an equal share of the random purchasers but no biased individuals. If the rival selects high complexity, the firm will share all of the random purchasers and simplicity biased equally, and compete for informed consumers.

$$\pi_i^R(L, p') = \left(\lambda_L \left[\frac{1-\mu_L}{2} + [1 - F^L(p')] \mu_L \right] + \lambda_H \left[\delta(1-\mu) + \frac{(1-\delta)(1-\mu)}{2} \right] \right) p'$$

When the firm selects low complexity under reversed supports: If the rival selects low, the random purchasers and simplicity biased are shared equally. The firms compete for informed buyers. If the rival selects high complexity, the firm wins all of the simplicity biased and an equal share of random purchasers but attracts no informed consumers.

The following two payoffs are for the payoffs of a firm that deviates to a price from the support of the other distribution, maintaining the reversed support.

$$\pi_i^R(H, p') = \left(\lambda_L \left[[1 - F^L(p')] \mu + \frac{(1 - \delta)(1 - \mu)}{2} \right] + \lambda_H \left[\frac{(1 - \mu_H)}{2} \right] \right) p'$$

If the rival selects low complexity, the firm wins an equal share of the random purchasers, competes for informed consumers and does not attract any biased consumers.

$$\pi_i^R(L, p) = \left(\lambda_L \left[\frac{1 + \mu_L}{2} \right] + \lambda_H \left[\delta(1 - \mu) + \frac{(1 - \delta)(1 - \mu)}{2} + \mu[1 - F^H(p)] \right] \right) p$$

If the rival chooses low, the firm shares the random purchasers and simplicity biased equally. The firm also attracts all informed consumers because the firm has selected a price from the support of $F^H(p)$, which under reversed supports must be a lower price than the rival. If the rival selects high complexity, the firm shares the random purchasers, wins all of the simplicity biased and competes for informed consumers.

Derivation of Firm Profit

Firm profit is given by (4.1). Rearrange to find:

$$\pi_i = \frac{1}{2} \cdot [\lambda_L \cdot (\delta\mu - \delta + \mu_H - \mu) + 1 - \mu_H]$$

Insert λ_L from (4.11).

$$\begin{aligned} \pi_i &= \frac{1}{2} \cdot \left[\frac{\mu - \mu_H + \delta - \delta\mu}{\mu_L - \mu_H} \cdot (\delta\mu - \delta + \mu_H - \mu) + 1 - \mu_H \right] \\ \pi_i &= \frac{1}{2} \cdot \left[1 - \mu_H - \frac{(\mu - \mu_H + \delta - \delta\mu)^2}{\mu_L - \mu_H} \right] \end{aligned}$$

This provides a closed-form solution for firm and industry profit, defined entirely by $(\delta, \mu_L, \mu, \mu_H)$. ■

Derivation for $F^L(p')$ when $\lambda_L = 1$:

Note that when $\lambda_L = 1$ exactly, $\delta = \frac{\mu_L - \mu}{1 - \mu}$. Therefore, in the expression for $F^L(p')$ given in (4.16), the right hand term is zero. Rearrange to yield $\mu = \frac{1 - \mu_L}{1 - \delta}$. Insert this into the expression for $F^L(p')$ to yield:

$$F^L(p) = \frac{1}{2} + \frac{1}{2\mu_L} \left(1 - \frac{1 - \mu_L}{p'} \right)$$

Add and subtract $\frac{1}{2}$.

$$F^L(p) = 1 + \frac{1}{2\mu_L} \left(1 - \frac{1 - \mu_L}{p'} \right) - \frac{1}{2}$$

$$F^L(p) = 1 + \frac{1}{2\mu_L} - \frac{1 - \mu_L}{2\mu_L p'} - \frac{1}{2}$$

$$F^L(p) = 1 + \frac{1 - \mu_L}{2\mu_L} - \frac{1 - \mu_L}{2\mu_L p'}$$

$$F^L(p) = 1 + \frac{1 - \mu_L}{2\mu_L} \left(1 - \frac{1}{p'} \right)$$

$$F^L(p) = 1 - \frac{1 - \mu_L}{2\mu_L} \left(\frac{1 - p'}{p'} \right)$$

Therefore, when $\lambda_L = 1$ exactly, the price distribution collapses. ■

Proof of Optimal Obfuscation with a Decreasing Marginal Effect of Complexity on Consumer Sophistication

To show that the main equilibria are preserved we follow the same structure as the proof of Lemma 1 & Proposition 1.

Step 1: Establish that when $\mu_L - \mu > \mu - \mu_H$ and $\delta < \frac{\mu - \mu_H}{1 - \mu}$, the element $(H, 1)$ continues to be a component of the equilibrium mixed strategy. This requires:

$$\lambda_L \cdot (1 - \delta)(1 - \mu) + \lambda_H \cdot (1 - \mu_H) \geq \lambda_L \cdot (1 - \mu_L) + \lambda_H \cdot (1 + \delta)(1 - \mu)$$

$$\lambda_L \cdot (\mu_L - \mu - \delta + \delta\mu) + \lambda_H \cdot (\delta\mu - \delta + \mu - \mu_H) \geq 0$$

The condition is satisfied for any $\delta \in [0, \frac{\mu - \mu_H}{1 - \mu}]$. Note that $\delta < \frac{\mu - \mu_H}{1 - \mu}$ implies $\delta < \frac{\mu_L - \mu}{1 - \mu}$ because of the decreasing marginal effect of complexity: $\mu_L - \mu > \mu - \mu_H$.

Step 2: Replicate the proof of the optimal obfuscation conditions in Proposition 1, which do not hinge on the constant marginal effect of obfuscation assumption. The derivations for the main model continue to hold. ■

Chapter 5

Conclusion & Further Research

The objective of this thesis was to provide a comprehensive examination and contribution to the intersecting literatures of price competition and strategic obfuscation. The overarching purpose has been to establish how firms' obfuscation and pricing incentives are influenced by consumers' distaste for complex pricing and complicated products. This sharpens our understanding of the behaviour of competing firms, generates clear guidelines for regulatory policy and provides testable predictions for experimental and empirical research.

In this concluding statement, I review the main contributions of each of the three main chapters of this thesis and synthesise the central themes of these research projects. I also discuss the limitations of both the research contained in this thesis and the broader research area. Finally, I offer several avenues for future research that have developed directly from the production of this thesis.

5.1 Conclusions

The purpose and contribution of Chapter 2 was to provide an authoritative account of the recently developed literatures of strategic obfuscation and price competition. I present a unified treatment of the most important developments in this established body of research. Moreover, Chapter 2 highlights several inconsistencies between the assumptions and predictions of theoretical models, and the results reported in experimental and empirical studies.

The review article also emphasises that mainstream theoretical models frequently overlook how consumers might respond to the obfuscation that they face in retail markets. More specifically, mainstream literature tends to build upon the benchmark model of sales, due to Varian (1980), where consumers who cannot compare prices choose a seller at random. Section 2.4 specifically investigates how consumers might respond to obfuscation strategies in practice. This section emphasises empirical and experimental evidence that finds consumers actively discriminate against obfuscating firms. Previously, Gaudeul and Sugden (2012) and Chioveanu and Zhou

(2013) were the only two studies to incorporate such behaviour in the Industrial Economics literature.

The research contained in this thesis develops these ideas further by explicitly studying: The effects of simplicity biased consumers on the obfuscation strategies of sellers, the effects on firm profit and consumer welfare, and the implications for policy. This thesis has also explored the opportunities for firms to inflate and suppress consumers' biases, and the opportunities for firms to combine complex pricing with additional advertising campaigns to manipulate complexity-averse consumers, which generates a novel form of false advertising.

The review article also identifies that intervention by competition authorities has the potential to generate greater harm than good for market competition and consumers. In the obfuscation literature, this conclusion was primarily established in markets with asymmetric firms. Therefore, in the subsequent chapters of this thesis, the effects of policy measures designed to enhance consumer sophistication and limit the effects of obfuscation, are investigated in an environment with identical sellers. The symmetric analysis removes any results that are predicated on asymmetric market prominence or production costs, which are required for at least one firm to depart from selecting maximum obfuscation in several previous studies.

The contribution of Chapter 3 was to develop a comprehensive game theoretic model of competition, which incorporates and focuses on the effects of simplicity biased consumers. This sharpens our understanding of how consumers' aversion to obfuscation might affect market dynamics in practice. The model demonstrates that obfuscation becomes significantly less attractive to firms once we account for simplicity biases in consumers' decision-making. Moreover, it is possible that the presence of biased consumers prevents firms from obfuscating entirely. Therefore, biased consumers generate greater market transparency and may also reduce the prices that consumers pay.

This is especially interesting because the firm with the simplest pricing scheme is demonstrated to charge the highest price in equilibrium. Therefore, consumers who are biased towards simple pricing structures will pay a higher price than if they had purchased from a randomly chosen seller. However, their presence generates a positive externality on the market by tempering the ability of firms to obfuscate. This allows more consumers to compare prices and stimulates price competition.

Chapter 3 also investigates the implications of policies to improve consumer sophistication in the presence of simplicity biased consumers. This is motivated by the need for further research in this area, identified in the review article in Chapter 2. The model predicts that policies to increase the fraction of consumers that are able to compare prices for any given level of market complexity, will have the most significant effects in markets where: Consumer sophistication and engagement is low, obfuscation is high and consumers are strongly biased towards simple pricing. These types

of policies can also have a multiplied effect by simultaneously improving consumers' abilities to overcome obfuscation and reducing firms' incentive to obfuscate.

Finally, the model in Chapter 3 introduces a simple cost penalty for firms that obfuscate to reflect the idea that producing a complicated pricing scheme requires effort and may generate additional costs for the firm. As anticipated, obfuscation costs reduce the attractiveness of obfuscation. However, it is also discussed that regulatory policy that targets this direct cost of obfuscation can enhance the competitive process. This is particularly interesting in circumstances where policies to improve consumer sophistication generate adverse incentives for firms.

In the model presented in Chapter 3, firms select their price and obfuscation level at the same time as their rival. However, both sellers commit to an obfuscation level before choosing their prices. This reflects the established idea that prices can be adjusted quickly but obfuscation cannot. Therefore, the starting point for Chapter 4 is to consider whether the effects of simplicity biased consumers are preserved when pricing and obfuscation decisions are taken simultaneously. The rationale follows that in some environments, obfuscation may be easily adjustable.

To study these effects, the model analysed in Chapter 3 is restated and significantly extended. It is demonstrated that simplicity biases by consumers continue to improve market transparency by suppressing firms' incentives to obfuscate. However, in contrast with the sequential model, the seller with the lowest price will now choose a lower level of obfuscation than the less competitive rival because the intuition for obfuscation differs. In the sequential framework in Chapter 3, a firm with a lower level of obfuscation attracts a larger share of the confused population, which weakens the incentive for this firm to compete in prices for informed consumers. In the simultaneous framework in Chapter 4, a firm with a low price maximises profit when there exists a higher proportion of informed consumers than a firm with a high price, which generates the opposite correlation.

Further contributions developed in Chapter 4 include the consideration of endogenous simplicity biases, which to the best of my knowledge have not been studied before in the context of obfuscation. This allows the degree of simplicity bias to be influenced by a firm's marketing efforts and regulatory policy. It is demonstrated that firms will engage in a novel form of false advertising whereby sellers select a complex pricing scheme to limit comparability with other firms, but choose a simple advertising campaign that persuades consumers that their pricing structure is simple.

These predictions are consistent with recent developments in retail pricing across the groceries, insurance, finance and telecommunications sectors and provides a new area of concern for regulators. Chapter 4 also shows that firms will dissuade consumers from purchasing on the basis of simplicity. This allows firms to increase their obfuscation and charge higher prices.

The models presented in this thesis share several common predictions that are consistent with, and develop, the existing literature. Firstly, firms will randomise their prices according to a mixed strategy equilibrium. Secondly, firms tend to earn higher profits when they are able to obfuscate their prices. Thirdly, policies to remedy obfuscation can both aid and harm consumers. Fourthly, the introduction of simplicity biased consumers suppresses firms' incentives to obfuscate. However, several of these results are surprising. This includes the welfare prediction contained in Chapter 3: Simplicity biased consumers pay the highest price but their presence can reduce average prices. It is also surprising that firms will attempt to simultaneously use complex pricing and attempt to convince consumers that their pricing scheme is less complex than their rivals.

In each chapter, the limitations of both the theoretical assumptions and modelling approaches have been evaluated. However, it is instructive to report several broader limitations of this theoretical area of research that are highlighted by the projects contained in this thesis. Firstly, the assumptions underpinning consumer behaviour in theoretical models play a critical role in determining both the optimal strategies of sellers and the effects of policy measures. It has been demonstrated throughout this thesis that neglecting consumers' aversion to complex pricing will lead industrial economists to overestimate the strengths of firms' incentives to obfuscate their prices. This overstates the motivation for intervention by regulators. Therefore, the hypothesised behaviour of demand-side agents must reflect the behaviour observed in experimental and empirical studies, to ensure that the predictions and policy implications are attuned to real world market dynamics.

Secondly, this thesis reinforces that the predictions of theoretical models are inextricably linked to the timing of the game that agents play. Therefore, it is critical to motivate the timing structure from observed empirical and experimental evidence. Throughout this thesis, I also recognise that obfuscation can follow a wide variety of forms. The models developed in these research projects are intentionally general and primarily follow an interpretation of complex pricing schemes, although many other interpretations are equally plausible.

Thirdly, the academic interest in price obfuscation focuses almost universally on retail markets, in which sophisticated firms interact with imperfectly informed consumers. There are two main caveats with this approach. Firstly, obfuscation is frequently observed in commercial markets where both the seller and buyer are informationally sophisticated agents. Therefore, it is not clear that the explanations for sellers to obfuscate in retail markets, based on cognitive limitations, information deficiencies and inattention of buyers, can explain the persistence of obfuscation in commercial settings. This is discussed in the subsequent future research avenues in Section 5.2.

Secondly, this literature examines the behaviour of two types of agents: Buyers and sellers. However, in most modern marketplaces there exists intermediaries. This might refer to a hired advisor to support a buyer choosing between products or a price comparison website that aids

consumers' ability to compare the prices available. This introduces an additional layer of incentives within the market that has not been thoroughly analysed in a framework of competitive obfuscation. For example, obfuscation by sellers might stimulate consumers' demand for intermediary services. This line of thought is particularly interesting in markets where one firm owns both a competing seller in the market and a market intermediary, which is one additional avenue for further research.

The models presented within these articles and across this strand of literature also frequently assume that product differentiation is negligible and firms are equally positioned, in terms of market prominence and consumer loyalty. These models draw closer to reality in markets where product differentiation is being eroded. However, caution should be exercised when applying theoretical frameworks grounded in an environment with symmetric competitors to markets where there exists clear market leadership effects and heterogeneity in the level of consumer loyalty retained by sellers. This thesis has also not considered the interesting strategies used by firms to stimulate consumer loyalty, which is an interesting area for future research within the pricing and obfuscation literatures. The assumption of unit demand from buyers is also prevalent across the related literatures of obfuscation and price competition. In practice, the quantity that each consumer will buy tends to be a function of both the price of the product and the hassle that consumers face in the market. Incorporating less rigid forms of demand is an opportunity for future research.

5.2 Further Research Projects

This thesis is completed by outlining two future research projects derived specifically from the work presented in this thesis.

5.2.1 Low Income Consumers and Add-On Insurance Products

In Ellison's (2005) seminal model of add-on pricing, firms sell a base good and an add-on consumable to two types of consumers: High and Low. Low types are low income consumers with a higher marginal utility of money. They are assumed to be more sensitive to price differences across sellers and less likely to purchase higher quality upgrades than their high income counterpart.

Ellison (2005) shows that there exists a lower bound on the price that a firm will charge for the base product. The intuition follows that if a firm selects a price below its competitors, more consumers will buy from the firm. However, the additional demand from consumers who switch to the firm corresponds primarily to low types, who cannot afford to purchase the profitable add-on. Therefore, the existence of these 'cheapskates' restricts price competition in the market for the base product.

The assumption that high income individuals hold a higher propensity for add-ons fits well in many markets. For example, higher income individuals may be more willing to buy the manufacturer's matching carry case for a new laptop or a dealership servicing package for a new car. However, in some cases this assumption is tenuous. For example, when purchasing a flight ticket, consumers have the option to buy additional insurance that allows them to cancel or rearrange their flight without additional charges. Neither of the high or low income consumers anticipate cancelling their flight but the possibility of additional charges may weigh more heavily on the low income individual. Furthermore, it may be reasonable to assume that higher income individuals fly more frequently. Therefore, insurance may be less worthwhile to the high income consumer because the possibility of additional charges for unexpected events on a single occasion are distributed over a higher quantity of total flights. This connects intimately to the topic of risk aversion across consumers with heterogeneous incomes.

In this situation, low income individuals who are more sensitive to price differences between sellers, are the main consumers of add-ons. This may challenge the main mechanism in Ellison's (2005) model in which base products feature low mark-ups and add-ons are profit generators. For example, will the add-on remain the profitable item or will the base good become the profitable product? Will firms instead compete on the combined price of the base and add-on product? Can the price of the base product exceed the combined price for both products?

5.2.2 Commercial Obfuscation

Obfuscation adopts a wide variety of forms. However, economic analysis has been almost exclusively confined to retail markets where sophisticated firms interact with imperfectly informed consumers. However, in commercial markets, both the buyer and seller are informationally sophisticated. It might, therefore, appear that obfuscation can play no role because buyers are engaged and able to compare the prices available from other suppliers. In practice, however, obfuscation is common in commercial markets.

It would, therefore, be interesting to conduct a joint theoretical and empirical analysis to understand the incentives for firms selling to commercial buyers to obfuscate their prices. This would include an analysis of the commercial markets most saturated by obfuscation, the different forms that obfuscation might follow and the methods used by buyers to overcome obfuscation in a commercial setting. This direction of research may also yield novel areas of intervention by regulators and competition authorities outside of their primary remit of retail markets.

Bibliography

- Acutt, M., Elliott, C. and Robinson, T. (2001), 'Credible regulatory threats', *Energy Policy* **29**(11), 911–916.
- Agnew, J. R. and Szykman, L. R. (2005), 'Asset allocation and information overload: The influence of information display, asset choice, and investor experience', *Journal of Behavioral Finance* **6**(2), 57–70.
- Akman, P. and Garrod, L. (2011), 'When are excessive prices unfair?', *Journal of Competition Law and Economics* **7**(2), 403–426.
- Armstrong, M., Vickers, J. and Zhou, J. (2009), 'Consumer protection and the incentive to become informed', *Journal of the European Economic Association* **7**(2-3), 399–410.
- Bachi, B. and Spiegel, R. (2015), 'Buridanic competition', *Working Paper* .
- Bakos, J. Y. (1997), 'Reducing buyer search costs: Implications for electronic marketplaces', *Management Science* **43**(12), 1676–1692.
- Bertrand, M., Karlan, D., Mullainathan, S., Shafir, E. and Zinman, J. (2010), 'What's advertising content worth? Evidence from a consumer credit marketing field experiment', *The Quarterly Journal of Economics* **125**(1), 263–306.
- Bordalo, P., Gennaioli, N. and Shleifer, A. (2013), 'Salience and consumer choice', *Journal of Political Economy* **121**(5), 803–843.
- Bordalo, P., Gennaioli, N. and Shleifer, A. (2016), 'Competition for attention', *The Review of Economic Studies* **83**(2), 481–513.
- Brown, J., Hossain, T. and Morgan, J. (2010), 'Shrouded attributes and information suppression: Evidence from the field', *The Quarterly Journal of Economics* **125**(2), 859–876.
- Bruce, A. C. and Johnson, J. E. (1996), 'Decision-making under risk: Effect of complexity on performance', *Psychological Reports* **79**, 67–76.

- Campbell, M. C. (2007), "'Says who?' How the source of price information and affect influence perceived price (un)fairness', *Journal of Marketing Research* **44**(2), 261–271.
- Carlin, B. I. (2009), 'Strategic price complexity in retail financial markets', *Journal of Financial Economics* **91**(3), 278–287.
- Carlin, B. I. and Manso, G. (2011), 'Obfuscation, learning, and the evolution of investor sophistication', *Review of Financial Studies* **24**(3), 754–785.
- Chioveanu, I. and Zhou, J. (2013), 'Price competition with consumer confusion', *Management Science* **59**(11), 2450–2469.
- Conlisk, J., Gerstner, E. and Sobel, J. (1984), 'Cyclic pricing by a durable goods monopolist', *The Quarterly Journal of Economics* **99**(3), 489–505.
- Dahremoller, C. (2013), 'Unshrouding for competitive advantage', *Journal of Economics & Management Strategy* **22**(3), 551–568.
- De Clippel, G., Eliaz, K. and Rozen, K. (2014), 'Competing for consumer inattention', *Journal of Political Economy* **122**(6), 1203–1234.
- De los Santos, B., Hortacsu, A. and Wildenbeest, M. R. (2012), 'Testing models of consumer search using data on web browsing and purchasing behavior', *American Economic Review* **102**(6), 2955–2980.
- Dertwinkel-Kalt, M., Köhler, K., Lange, M. and Wenzel, T. (2016), 'Demand shifts due to salience effects: Experimental evidence', *Journal of the European Economic Association* (Forthcoming).
- Eliaz, K. and Spiegler, R. (2011a), 'Consideration sets and competitive marketing', *The Review of Economic Studies* **78**(1), 235–262.
- Eliaz, K. and Spiegler, R. (2011b), 'On the strategic use of attention grabbers', *Theoretical Economics* **6**(1), 127–155.
- Ellison, G. (2005), 'A model of add-on pricing', *The Quarterly Journal of Economics* **120**(2), 585–637.
- Ellison, G. (2006), 'Bounded rationality in industrial organization', *MIT and NBER* (Available at: <http://economics.mit.edu/files/904>. Accessed January 2016.).
- Ellison, G. and Ellison, S. F. (2009), 'Search, obfuscation, and price elasticities on the internet', *Econometrica* **77**(2), 427–452.
- Ellison, G. and Wolitzky, A. (2012), 'A search cost model of obfuscation', *The Rand Journal of Economics* **43**(3), 417–441.

- Frank, R. G. and Lamiraud, K. (2009), 'Choice, price competition and complexity in markets for health insurance', *Journal of Economic Behavior and Organization* **71**(2), 550–562.
- Fudenberg, D. (2006), 'Advancing beyond "advances in behavioral economics"', *Journal of Economic Literature* **44**(3), 694–711.
- Gabaix, X. and Laibson, D. (2006), 'Shrouded attributes, consumer myopia, and information suppression in competitive markets', *The Quarterly Journal of Economics* **121**(2), 505–540.
- Gaudeul, A. and Sugden, R. (2012), 'Spurious complexity and common standards in markets for consumer goods', *Economica* **79**(314), 209–225.
- Grubb, M. D. (2015a), 'Failing to choose the best price: Theory, evidence, and policy', *Review of Industrial Organization* **47**(3), 303–340.
- Grubb, M. D. (2015b), 'Overconfident consumers in the marketplace', *Journal of Economic Perspectives* **29**(4), 9–36.
- Gu, Y. and Wenzel, T. (2014), 'Strategic obfuscation and consumer protection policy', *The Journal of Industrial Economics* **62**(4), 632–660.
- Gu, Y. and Wenzel, T. (2015), 'Putting on a tight leash and levelling playing field: An experiment in strategic obfuscation and consumer protection', *International Journal of Industrial Organization* **42**, 120–128.
- Haan, M. A. and Moraga-González, J. L. (2011), 'Advertising for attention in a consumer search model', *The Economic Journal* **121**(552), 552–579.
- Hastings, J. and Tejada-Ashton, L. (2008), 'Financial literacy, information, and demand elasticity: Survey and experimental evidence from Mexico', *NBER Working Paper Series* **14538**.
- Homburg, C., Totzek, D. and Kramer, M. (2014), 'How price complexity takes its toll: The neglected role of a simplicity bias and fairness in price evaluations', *Journal of Business Research* **67**(6), 1114–1122.
- Hossain, T. and Morgan, J. (2006), '...plus shipping and handling: Revenue (non) equivalence in field experiments on ebay', *Natural Field Experiments* **6**(2).
- Huck, S. and Wallace, B. (2015), 'The impact of price frames on consumer decision making: Experimental evidence.', *Working Paper. Available at www.ucl.ac.uk/uctpbwa/papers/price-framing.pdf*

- Huck, S. and Weizsäcker, G. (1999), 'Risk, complexity, and deviations from expected-value maximization: Results of a lottery choice experiment', *Journal of Economic Psychology* **20**(6), 699–715.
- Ireland, N. J. (2007), 'Posting multiple prices to reduce the effectiveness of consumer price search', *The Journal of Industrial Economics* **55**(2), 235–263.
- Iyengar, S. S. and Kamenica, E. (2010), 'Choice proliferation, simplicity seeking, and asset allocation', *Journal of Public Economics* **94**(7), 530–539.
- Kahneman, D. and Tversky, A. (1979), 'Prospect theory: An analysis of decision under risk', *Econometrica* **47**(2), 263–292.
- Kalayci, K. (2015), 'Price complexity and buyer confusion in markets', *Journal of Economic Behavior and Organization* **111**, 154–168.
- Kalayci, K. and Potters, J. (2011), 'Buyer confusion and market prices', *International Journal of Industrial Organization* **29**, 14–22.
- Kim, H. M. and Kachersky, L. (2006), 'Dimensions of price salience: a conceptual framework for perceptions of multi-dimensional prices', *Journal of Product and Brand Management* **15**(2), 139–147.
- Kőszegi, B. (2014), 'Behavioral contract theory', *Journal of Economic Literature* **52**(4), 1075–1118.
- Kőszegi, B. and Rabin, M. (2006), 'A model of reference-dependent preferences', *The Quarterly Journal of Economics* **121**(4), 1133–1165.
- Kőszegi, B. and Szeidl, A. (2013), 'A model of focusing in economic choice', *Quarterly Journal of Economics* **128**(1), 53–104.
- Krishna, A., Briesch, R., Lehmann, D. R. and Yuan, H. (2002), 'A meta-analysis of the impact of price presentation on perceived savings', *Journal of Retailing* **78**(2), 101–118.
- Mador, G., Sonsino, D. and Benzion, U. (2000), 'On complexity and lotteries' evaluation - Three experimental observations', *Journal of Economic Psychology* **21**(6), 625–637.
- McGovern, G. and Moon, Y. (2007), 'Companies and the customers who hate them', *Harvard business review* **85**(6), 78–84.
- Miravete, E. J. (2006), 'The limited gains from complex tariffs', *Working Paper* .
- Miravete, E. J. (2013), 'Competition and the use of foggy pricing', *American Economic Journal: Microeconomics* **5**(1), 194–216.

- Morgan, J., Orzen, H. and Sefton, M. (2006), 'An experimental study of price dispersion', *Games and Economic Behavior* **54**, 134–158.
- Muir, D., Seim, K. and Vitorino, M. A. (2013), 'Drip pricing when consumers have limited foresight: Evidence from driving school fees', *Working Paper* .
- Narasimhan, C. (1988), 'Competitive promotional strategies', *The Journal of Business* **61**(4), 427–449.
- Papi, M. (2014), 'Firm competition with tradeoff-avoiding consumers', *Working Paper* .
- Petrikaite, V. (2016), 'Consumer obfuscation by a multiproduct firm', *Working Paper* .
- Piccione, M. and Spiegler, R. (2012), 'Price competition under limited comparability', *The Quarterly Journal of Economics* **127**(1), 97–135.
- Prelec, D. and Loewenstein, G. (1991), 'Decision making over time and under uncertainty: A common approach', *Management Science* **37**(7), 770.
- Rhodes, A. (2015), 'Multiproduct retailing', *Review of Economic Studies* **82**(1), 360–390.
- Rhodes, A. and Wilson, C. (2015), 'False advertising', *Working paper available at <http://dx.doi.org/10.2139/ssrn.2716692>* .
- Salant, Y. and Rubinstein, A. (2008), '(a, f): Choice with frames', *The Review of Economic Studies* **75**(4), 1287–1296.
- Salant, Y. and Siegel, R. (2015), 'Contracts with framing', *Working Paper* .
- Shelegia, S. (2012), 'Multiproduct pricing in oligopoly', *International Journal of Industrial Organization* **30**(2), 231–242.
- Sitzia, S. and Zizzo, D. J. (2011), 'Does product complexity matter for competition in experimental retail markets?', *Theory and Decision: An International Journal for Methods and Models in the Social and Decision Sciences* **70**(1), 65–82.
- Sobel, J. (1984), 'The timing of sales', *The Review of Economic Studies* **5**(3), 353–368.
- Sonsino, D., Benzion, U. and Mador, G. (2002), 'The complexity effects on choice with uncertainty - Experimental evidence', *The Economic Journal* **112**(482), 936–965.
- Spiegler, R. (2006), 'Competition over agents with boundedly rational expectations', *Theoretical Economics* **1**(2), 207–231.
- Spiegler, R. (2011), *Bounded Rationality and Industrial Organization*, New York ; Oxford University Press.

- Spiegler, R. (2014), 'Competitive framing', *American Economic Journal: Microeconomics* **6**(3), 35–58.
- Stahl, D. O. (1989), 'Oligopolistic pricing with sequential consumer search', *The American Economic Review* **79**(4), 700.
- Tirole, J. (1988), *The Theory of Industrial Organization*, Cambridge (Mass.): MIT Press.
- Varian, H. R. (1980), 'A model of sales', *The American Economic Review* **70**(4), 651–659.
- Vives, X. (2001), *Oligopoly Pricing*, MIT Press.
- Wenzel, T. (2014), 'Consumer myopia, competition and the incentives to unshroud add-on information', *Journal of Economic Behavior and Organization* **98**, 89–96.
- Wilson, C. M. (2010), 'Ordered search and equilibrium obfuscation', *International Journal of Industrial Organization* **28**(5), 496–506.
- Wilson, C. M. and Waddams-Price, C. (2010), 'Do consumers switch to the best supplier?', *Oxford Economic Papers* **62**(4), 647–668.
- Xia, L., Monroe, K. B. and Cox, J. L. (2004), 'The price is unfair! A conceptual framework of price fairness perceptions', *Journal of Marketing* **68**(4), 1–15.
- Zhou, J. D. (2014), 'Multiproduct search and the joint search effect', *American Economic Review* **104**(9), 2918–2939.