ADD-ON MARKETS WITH NAÏVE CONSUMERS

Ioannis Pappous

Thesis submitted for the degree of
Doctor of Philosophy

University of East Anglia
School of Economics

January 2019

Abstract

This thesis aims to improve our understanding of firms' incentives to offer add-ons, as well as the implications that selling add-ons has for competition and welfare. It consists of three studies.

In the first study we develop two models that explain the establishment of add-on pricing equilibria in markets with physical add-on services (e.g. hotels, airlines), and in software markets (e.g. apps with microtransactions). We show that add-on pricing equilibria can emerge despite consumers being fully rational and despite the presence of a binding price floor.

The second and third studies analyse markets with naïve consumers. The goal of both studies is to explore the extent to which the incentive to exploit naïve consumers' mistakes affects firms' choices of horizontal product differentiation in the core product market.

In the second study we analyse a model in which add-ons are unwanted – they are never purchased by sophisticated consumers. We identify conditions under which firms offer maximally differentiated products in order to soften competition, and conditions under which each firm offers a product similar to its rival's in order to maximise profits from exploitation.

In the third study we develop a model in which add-ons are desirable – sophisticated consumers buy them under certain circumstances. We find that the sophisticates' demand for add-ons complicates the relationship between prices and the degree of horizontal product differentiation. This raises the possibility that intermediate product differentiation also emerges in equilibrium.

Both the second and the third study show that benefits to consumers from the firms' choices of product differentiation can more than offset the harm to consumer welfare from the exploitation of consumer mistakes. Thus, whether conventional regulatory interventions have a positive effect on consumer welfare depends, to a large extent, on their effect on firms' product differentiation incentives.

Contents

List of Figures

Acknowledgements

Working towards a PhD was a challenging and rewarding experience, through which I grew as a researcher, an economist, and a person. This journey would not be the same without all the people that shaped it.

I am greatly indebted to my supervisors, Prof. Bruce Lyons and Prof. Kai-Uwe Kühn, for their invaluable guidance and continuous encouragement. I am deeply grateful for all the time they dedicated to discussing ideas, responding to my many questions and reviewing my work.

My sincere thanks also go to my former supervisor, Dr. Sang-Hyun Kim, for so patiently introducing me to the world of academic research.

My thanks are extended to all the faculty at the School of Economics, the Centre for Competition Policy, and the participants of the IO reading group for fruitful discussions and for providing an excellent learning environment. I especially want to thank Prof. Stephen Davies, Prof. Robert Sugden, Prof. Mich Tvede, Dr. Farasat Bokhari, and Dr. Franco Mariuzzo for showing genuine interest in my study and for providing insightful advice.

I also wish to thank my fellow PhD students and office mates, as well as the numerous friends that I made over the past five years in Norwich, for the good times and countless laughs that we shared.

Finally, I wish to thank my family for their endless love, faith, and support. I am forever grateful to them for being there for me, and I owe to them everything that I have achieved.

We shall not cease from exploration,
and the end of all our exploring
will be to arrive where we started
and know the place for the first time.

T.S. Elliot

Chapter 1

Introduction

1.1 Motivation

In a wide range of industries, firms offer add-ons alongside their core services. Traditional examples include hotels selling breakfast and access to a minibar, airlines selling on-board meals and extra luggage space, and durable goods sellers offering extended warranties and repair services. The practice of separating core from add-on services has become particularly common for sellers of digital goods. Many smartphone apps, for instance, offer functionality and aesthetic upgrades via in-app purchases. Finally, the sale of add-on services is widely used in retail financial services. For example, holders of personal current accounts can pay for access to an overdraft service, while credit card users have the option of delaying repayments at the cost of interest and late payment charges.

This thesis aims to improve our understanding of firms' incentives to offer add-ons, as well as the implications that selling add-ons has for competition and welfare. What market characteristics drive firms to separately sell core and add-on services? How does the sale of add-ons affect consumer prices? How does it affect firm behaviour in other dimensions? What is the role of boundedly rational consumers? This thesis seeks to shed some light on these questions.

We define *add-on pricing* as the business strategy in which a firm separately sells a core service, or *base good*, and an add-on (component or service). We define *all-inclusive pricing* as the business strategy in which a firm sells only a bundle that includes the base good and the add-on. Any add-on is usually characterized by three features that distinguish the study of add-on pricing from the study of related but distinct problems. First, the add-on cannot be purchased separate from the base good. For example, a consumer cannot buy an on-board meal without a flight ticket, or use an overdraft service without a personal current account. This feature distinguishes the study of add-on pricing from the study of more traditional bundling and tying problems in markets with two products that each can be used to some capacity on its own (e.g. Whinston, 1990; Chen, 1997).

Second, inherent switching costs allow add-on pricing firms to enjoy some market power over the add-on market. For example, a flight passenger cannot buy on-board meals from a different supplier; a credit card user can secure funds from other sources, but that often carries inconvenience costs. Because of this market power, add-on pricing models feature similar characteristics with models in the switching cost literature, most notable being the subsidised first-period prices (Klemperer, 1987). The inherent nature of such switching costs distinguishes the study of add-on markets from the study of markets in which firms may choose to actively monopolise secondary markets in order to raise prices (Shapiro, 1995; Borenstein *et al.* 2000) or exclude rivals (Carlton and Waldman, 2002; Kühn and van Reenen, 2009)

Third, the purchase of add-ons is optional and independent of the decision to buy the base good. For instance, a bank account holder may choose to reduce spending rather than go into overdraft, while an app user may prefer to be interrupted by advertisement instead of paying for the upgraded, ad-free version of the same app. This distinguishes add-on pricing from partitioned pricing (e.g. Office of Fair Trading, 2013; Greenleaf *et al.*, 2016), in which mandatory additional surcharges are revealed during the transaction.[1]

[1] A related term, *drip pricing*, describes a business strategy in which firms reveals various charges at the later stages of a transaction. Drip pricing may be closer to add-on or partitioned pricing, depending on whether the dripped price is optional (e.g. purchase of one-day delivery) or mandatory (e.g. booking fees). For more information on drip pricing see Shelanski *et al.* (2012).

The thesis consists of three theoretical essays, each with relevant policy implications. The first essay, in Chapter 2, is motivated by the popularity of add-on pricing strategies in the real world vis-à-vis a theoretical result that add-on pricing is, in many cases, not an equilibrium strategy. More specifically, an influential paper by Ellison (2005) shows that firms which adopt an add-on pricing strategy may have limited ability to offer competitive deals. This effect results in supra-competitive profits if firms jointly adopt add-on pricing, but also implies that add-on pricing is not adopted in equilibrium.[2]

In Chapter 2 we develop two models that extend Ellison's framework. Each model explains the establishment of add-on pricing equilibria in different settings, even when add-on pricing limits firms' ability to offer competitive deals, and even when consumers are fully rational. The first model takes into account various cost savings that are, in some cases, associated with add-on pricing, and specifies conditions under which the benefit of such cost savings can outweigh the harm to add-on pricing firms' ability to compete. The presence of such cost savings is likely to play a role in the choice of business strategy of firms in markets with physical add-on services, such as hotels and airlines.

The second model recognizes that a consumer's choice of whether to buy add-on may not depend only on her valuation of the add-on, but also on her valuation of the base good. The model approximates a market of experience goods, in which any given product may or may not match consumers' tastes. In such a market, an add-on pricing business strategy can permit uninformed consumers to learn about a product's match quality by using the basic good as a "trial version". We show that the conditioning of add-on purchases on correct matching allows add-on pricing firms to price discriminate between consumers that are informed ex ante about match quality, and consumers that are uninformed. The possibility of such price discrimination can, under certain parameter constellations, establish an add-on pricing equilibrium even if add-on pricing constrains firms' ability to compete. Our model specification is particularly applicable to markets for digital products. For example, smartphone users can taste the

[2] We explain the model of Ellison (2005) in detail in the following section.

basic features of most apps for free and may buy optional functionality or cosmetic upgrades after good matching is realised.

Chapters 3 and 4 consider markets in which naïve consumers purchase add-on services by mistake. Prominent examples include retail financial markets, in which many consumers accidentally use overdraft facilities (Office of Fair Trading, 2008) or miss credit card payments (Gathergood *et al.*, 2018), and mobile telecommunications markets, in which many consumers accidentally exceed the call or data usage thresholds that are described in their contracts (Federal Communications Commission, 2010; Grubb, 2015b).

In such markets, firms have an incentive to alter their pricing strategies in order to profitably exploit naïve consumers' mistakes. Optimal pricing is described by relatively low core product prices – what we refer to as *upfront prices* – and relatively high add-on prices. The aim of both chapters is to explore the extent to which firms' incentive to profitably exploit naïve consumers affects non-price strategic variables. The strategic variable we study in both chapters is horizontal product differentiation. In studying this question, both chapters contribute to a growing literature that studies distortive effects of consumer exploitation.[3]

In Chapter 3 we model a market in which add-ons are unwanted – they are never bought by sophisticated consumers. We find that the degree of horizontal product differentiation affects prices in a non-linear way, and critically influences how much profit firms can retain from the sale of add-ons. When the choice of horizontal differentiation is endogenous, we identify conditions under which firms offer maximally differentiated products in order to soften competition, and conditions under which each firm prefers to offer a product similar to its rival's in an attempt to maximize add-on profits. We find the surprising result that the presence of naïve consumers may in fact lead to lower profits compared to the benchmark case in which all consumers are sophisticated. Conventional interventions, such as price controls and disclosure remedies may or may not improve consumer welfare.

In Chapter 4 we study how the findings of Chapter 3 change when add-ons are desirable – sophisticated consumers buy them under certain circumstances. We find

[3] See, for instance, Heidhues *et al.* (2016). We review this literature in greater detail in section 1.2.

that the sophisticated consumers' demand for add-ons results in a more complicated relationship between the degree of horizontal product differentiation and prices. Depending on how close substitutes the products are, firms may compete on upfront prices to attract both consumer types, on add-on prices in order to attract sophisticated types, or may choose add-on prices randomly as part of a mixed strategy. The choice of product differentiation becomes more complicated as well: products may feature maximum, minimum, or even intermediate product differentiation. This possible emergence of intermediate product differentiation is hard to explain using standard models. Conventional interventions may or may not be desirable both from a social and a consumer welfare perspective.

1.2 Issues in markets with add-ons: A review of the literature

A crucial question regarding markets with add-ons is whether firms can enjoy supra-competitive profits from add-on purchases. In section 1.2.1 we explain the profit irrelevance result that arises in most model settings and relate this result to loss-leader pricing and the literature on switching costs. We also discuss a version of the profit irrelevance result that emerges in markets with naïve consumers. Finally, we discuss literature in which firms can retain some extra profits from selling add-ons.

Add-on pricing is, at its core, a second-degree price discrimination strategy, aimed at extracting higher surplus from high-value consumers. Sections 1.2.2 and 1.2.3 study its role as such. In markets with rational consumers, add-on pricing allows firms to extract more surplus from consumers that value add-ons the most. The literature in section 1.2.2 studies firms' decision to sell add-ons separately or bundle them with base goods when such *preference-based price discrimination* is possible.[4]

In some markets consumers often naïvely underestimate their demand for certain product features when making a purchase decision. In such markets, add-on pricing

[4] Stole (2007) provides an excellent review of classical preference-based price discrimination models of imperfectly competitive markets.

can permit a second-degree *naïvete-based price discrimination*.[5] To answer whether naïvete-based price discrimination indeed takes place when add-ons are sold separately, a crucial question is whether firms prefer to educate naïve consumers about their future add-on demand by *unshrouding*, or whether firms prefer to keep add-on information *shrouded* in order to hold naïve consumers up. With regards to firms' ability to price discriminate, shrouding is similar to unbundling in that it aims to extract more surplus from a few valuable consumers, while unshrouding is similar to selling to everyone at an all-inclusive price. Section 1.2.3 reviews the literature that studies firms' incentives to shroud/unshroud.

Section 1.2.4 discusses externalities that arise between sophisticated and naïve consumers in market settings with add-on services. Section 1.2.5 discusses market distortions that often arise when naïve consumers underestimate their demand for add-ons. Section 1.2.6 discusses the effectiveness of regulations as a means to protect naïve consumers from exploitation.

1.2.1 How add-ons affect profits

A classic result in the add-on pricing literature is that firms compete away add-on profits by subsidising base good prices to even below marginal cost. Intuitively, if firms in a perfectly competitive industry earn profit π from the sale of add-ons, then the zero-profit condition requires that each firm makes losses equal to $-\pi$ in the first period. The same logic applies when markets are imperfectly competitive: firms compete away add-on profits and end up earning no more than permitted by their market power on the base good.

This competition-for-the-market logic is used by Shapiro (1995), who argues that durable goods sellers have an incentive to discount their equipment to capture aftermarket profits. In an earlier work, Lal and Matutes (1994) show that the profits that retailers enjoy from impulse goods (which can be loosely interpreted as the add-ons) are used to subsidise loss-leader goods that attract consumers in the store.

[5] For a study of first and third-degree naïvete-based price discrimination, see Heidhues and Koszegi (2017).

Verboven (1999) also finds that firms compete away add-on profits in a model in which heterogeneous consumers make a purchase choice based on base good prices. He finds evidence from the automobile industry that is consistent with this finding.

A competition-for-the-market argument appears also in the switching cost literature.[6] For instance, Klemperer (1987) shows that firms can earn monopolistic profits from consumers that are locked with that firm due to switching costs. However, firms compete more fiercely in the first period in order to attract and lock these consumers in, dissipating anticipated future profits.

Dissipation of add-on profits also takes place in markets in which boundedly rational consumers underestimate future add-on costs and make purchase decisions based on base good prices only. For instance, in Armstrong and Vickers (2012), fees paid for overdraft services by naïve consumers incentivise firms to compete more fiercely in the personal current accounts market, selling accounts at below-cost prices. In some sense, competition for the base good at least partially[7] protects consumers from exploitation – what Heidhues and Koszegi (2018) refer to as the "safety-in-markets" effect.

Ellison (2005) is the first to show that complete dissipation of add-on profits does not always emerge, as an adverse selection effect may constrain firms ability to cut base good prices. Such an adverse selection takes place if lower base good prices attract a disproportionate amount of consumers that do not buy add-ons.[8] Ellison and Ellison (2009) find evidence of such an adverse selection effect in a market for computer parts. The large number of retail financial products with zero base price (e.g. DellaVigna and Malmendier, 2004; Armstrong and Vickers, 2012) can be explained by the possibility that an adverse selection takes place at negative prices.

Some works that model markets with boundledly rational consumers adopt simpler variants of Ellison's adverse selection mechanism. In Heidhues *et al.* (2016; 2017), setting base good prices below some exogenously set "price floor" is

[6] Farrell and Klemperer (2007) provide an excellent survey of this literature.

[7] Protection is partial if sophisticated consumers buy the base good but not the add-on. See Armstrong (2015) for an analysis of such rip-off externalities.

[8] For a more detailed review of Ellison (2005) see section 1.2.2.

unprofitable, as it attracts a large number of arbitrageurs that enjoy the subsidised prices and avoid add-ons at a small inconvenience cost. Exogenous price floors are also assumed in an extension of Armstrong and Vickers (2012) with regards to personal current accounts, as well as in the working paper version of Grubb (2015b) with regards to mobile phone contracts.

Price floors can arise endogenously if base goods and add-ons are imperfect substitutes. This is the case in Michel (2017), who models a market in which the add-on is an extended warranty. Cutting base good prices below a certain threshold may be unprofitable if consumers find it less costly to replace a faulty product with a new one instead of buying a warranty. A similar mechanism takes place in a model of printer systems (which include ink) and cartridges by Miao (2010). If firms set cartridge prices too low, consumers may prefer to buy a new printer system instead of a new cartridge.

In reality, price floors may also arise due to market regulations (below-cost pricing may be viewed as prohibited predatory behaviour; regulation in the mutual funds industry may rule out "signing bonuses", as modelled in Heidhues *et al.* 2017), due to consumers' perception of high upfront prices as a signal for quality (Wolinsky, 1983), or if consumers are suspicious of extremely low prices, expecting a trap.

Models in Chapters 2, 3, and 4 of this thesis join the list of works that study market outcomes when price floors constrain competition on base prices. Chapter 2 specifies conditions under which an add-on pricing equilibrium exist in markets in which a price floor is present. Chapters 3 and 4 study how the presence of a price floor affects horizontal product differentiation in the market.

1.2.2 Selling add-ons to rational consumers

The seminal paper of Ellison (2005) highlights the effectiveness of add-on pricing as a price discrimination device in a competitive framework. In his model, two firms sell imperfect substitutes to consumers with high and low marginal utility of income. Each firm offers a high-quality and a low-quality version of its product. In the "standard price discrimination game" each firm advertises the prices of both versions of its product. In the "add-on pricing game" each firm only advertises the price of its low-quality product version; a consumer can observe the price of the high-quality version only by visiting that firm.

Ellison shows that firm profits at the standard price discrimination game are lower than profits at the add-on pricing game. The reason is that, in the add-on pricing game, firms can only compete on the price of the low-quality version. Because of consumer heterogeneity in their marginal utility of income, price cuts attract a disproportionate number of low-type consumers that do not buy the high-quality version. This adverse selection effect limits price-cutting, softening competition.

While more profitable, following an add-on pricing strategy (that is, keeping the price of the high-quality version hidden) is never an individually rational firm strategy when the choice to advertise prices is endogenous. The reason is precisely the adverse selection effect described above. Given a firm's choice to hide the price of its high-quality version, it is always profitable for a rival to advertise their own high-quality version and undercut in order to steal some high-type consumers.

Several works explore market conditions that facilitate the existence of add-on pricing equilibria versus all-inclusive equilibria when Ellison's adverse selection effect is not present. Geng and Shulman (2015), emphasise on the role of cost savings, often cited to be the rationale behind unbundling add-ons, and heterogeneity in consumers' price sensitivity. They show that the firms' choice of pricing strategy depends on add-on production costs, on the relative price sensitivity of consumers that have demand for add-ons to that of consumers that do not, as well as on the existence of consumers that consume the add-on if and only if it is bundled with the base good. They note the possibility that an add-on pricing equilibrium may emerge despite leading to lower profits than an all-inclusive equilibrium.

Lin (2017) studies incentives to unbundle base goods from add-ons in a model in which two firms are vertically differentiated. The high-quality firm has an incentive to sell the add-on separately in order to price discriminate. The low-quality firm faces two conflicting incentives: selling the add-on separately permits that firm to price discriminate, while bundling core service and add-on allows that firm to differentiate itself from its high-quality rival. In equilibrium, the high-quality firm always sells the add-on as optional. The low-quality firm may bundle the add-on with the core service, sell it as optional, or not offer it at all, depending on the add-on's production cost-to-value ratio.

Geng *et al.* (2018) show that an upstream monopolist's choice of business strategy depends on the distribution contract this firm has with a downstream platform. In their model, add-ons, when sold separately, can be purchased directly from the monopolist. Under an agency contract, this effect makes add-on pricing more profitable than all-inclusive pricing. This is because it allows the monopolist to enjoy the full add-on revenue, while permitting it to pass some of the revenue loss from the lower base price to the platform through the commission rate (what they call the *loss-sharing effect*). Under a wholesale contract, a loss-sharing effect does not arise. In that case, the monopolist prefers to bundle base good with add-on, as that permits it to sell add-ons even to consumers that do not value them.

Cui *et al.* (2018) show that a monopolist's choice to unbundle the add-on depends on whether that monopolist can price discriminate between high and low-value consumers on the base good. Under uniform pricing on the base good, add-on pricing is more profitable when high base good valuation is correlated with high add-on valuation, as it allows the monopolist to price discriminate. Under discriminatory pricing on the base good, add-on pricing emerges when both high- and low-type consumers' add-on valuation is relatively low. This is because the optional nature of the add-on component can be used to capture consumers of each type that do not value the add-on service.

In a related question, Gomes and Tirole (2018) explain why some markets feature high add-on prices, while in other markets add-ons are given away for free. In their model, consumers may choose not to buy the base good if add-on prices are too high. The authors show that this possibility deters a base good monopolist from setting high add-on prices, and may even incentivise the monopolist to offer add-ons below cost. The higher the mark-up on the base good, the weaker the incentive to set high add-on prices. Hold-ups (i.e. high add-on prices resulting from the firm's inability to commit to lower prices) can occur when consumers are uninformed about add-on prices or when they naïvely ignore them; however, the extent of the hold-up is limited by the concern for missed sales.

Fruchter *et al.* (2011) study the same question in an earlier work by analysing a model in which add-on valuation is independent of base good valuation. They show that, in most cases, a monopolist is better off setting a high add-on price in order to extract the surplus of consumers that value the add-on. If high add-on valuation is associated with low base good valuation, then the monopolist may have an incentive to offer complimentary add-ons. This is because the free add-on yields additional surplus to consumers with low base good valuation (the low types). The monopolist then can charge a higher base price and extract more surplus from all consumers without driving away these low types.

The primary contribution of Chapter 2 relates to this literature, specifying conditions under which firms unbundle and sell add-ons separately from base goods even when Ellison's adverse selection effect takes place.

1.2.3 Shrouding add-ons (and other product features) from naïve consumers

The seminal paper of Gabaix and Laibson (2006) sparks the literature that studies shrouding incentives. In their model, naïve consumer types expect they have no demand for add-ons. They choose from which firm to purchase based on base good prices only, and end up buying add-ons at any price below their willingness to pay once they commit to a base good purchase. Sophisticated types, on the other hand, are aware of their future add-on demand, and so make a purchase choice based on expected lifetime costs.

Shrouding add-ons permits firms to sell add-ons at high prices to naïve types. Sophisticates types, expecting to be held up, substitute away. If at least one firm unshrouds, a fraction of naïve types become educated and act as the sophisticated types. If the fraction of naïve types is high enough then both firms shroud in equilibrium. No firm has incentive to unshroud, as a firm can only profit from unshrouding by cutting prices enough to prevent educated naïve types from substituting away. This is unprofitable, as it results in lower revenues from consumers that remain naïve.

If, instead, the fraction of naïve types is sufficiently low, each firm has an incentive to unshroud and sell add-ons to all consumers at lower prices. Shrouding is not profitable when rivals are unshrouding, as the only consumers that purchase add-ons from a shrouded firm are uneducated naïve types. Sophisticated and educated naïve types that purchase from a shrouding firm would expect to be held up and would substitute away.

Profits under both shrouding and unshrouding are independent of add-on sales, as firms compete away extra add-on profits by lowering base good prices.

Heidhues *et al.* (2017) is the first to study shrouding incentives when a binding price floor prevents complete dissipation of add-on profits. In their model, each firm sells a product that features an upfront price that is visible to all consumers and an unavoidable add-on fee that is hidden unless at least one firm unshrouds. Under shrouding, firms set the highest possible add-on fees, and compete upfront prices down to the price floor. Under unshrouding, firms compete over the total price, resulting in price competition a la Bertrand.

The authors show that, given no unshrouding costs, shrouding can only arise in equilibrium if the price combination under shrouding is higher than the intrinsic utility of the sold product. More specifically, the difference between price combination under shrouding and intrinsic utility must be high enough that each firm prefers to serve a fraction of the market at the shrouded prices, over unshrouding and serving the whole market at monopoly prices.

When products are socially valuable (i.e. their intrinsic utility is higher than their production costs), increasing the number of firms can bring about unshrouding in equilibrium. This is because profits under shrouding are shared between a higher number of firms, making deviations to unshrouding more profitable in comparison. On the contrary, when products are socially wasteful (i.e. their intrinsic utility is lower than their production costs) then shrouding is the unique equilibrium market outcome. This is because firms cannot profitably sell socially wasteful products at any price when unshrouding takes place.

The model in Chapter 2 of this thesis gives two alternative explanations as to why firms profitably sell add-ons to consumers when a price floor is present, without requiring the presence of boundedly rational consumers. The models in Chapters 3 and

4 take a shrouding equilibrium as a given (which is possible in Heidhues *et al.* (2017) even without above-monopoly prices if unshrouding costs are sufficiently high) in a market with naïve consumers that ignore add-on fees and sophisticated consumers that observe and take them into account in their purchase decision. The two models study second-order effects in firms' choice of horizontal differentiation under the assumption of such shrouding taking place.

A few works extend the model of Gabaix and Laibson (2006) to provide additional insights regarding shrouding incentives when add-on profits are fully competed away. Dahremoller (2013) assumes that the shrouding decision takes place before price-setting and that firms are heterogeneous in their add-on production costs. In this setting, unshrouding always emerges in a duopoly. Depending on parameter values, at least one firm always has an incentive to unshroud in order to reduce its rival's add-on profits, since lower add-on profits result in softer competition for the base good market.

Wenzel (2014) modifies the Gabaix and Laibson framework in a different dimension, studying how shrouding incentives may change depending on competition intensity. The author models a market in which the effect of unshrouding on the education of naïve types depends on the number of unshrouding firms. Each unshrouding firm reduces the number of naïve types in the market, thus incentivising rivals to unshroud as well. As a result, the parameter range for which unshrouding emerges increases with the number of firms in the market. For parameter values for which both shrouding and unshrouding equilibria exist, unshrouding dominates in terms of risk. This is because shrouding firms' add-on revenues fall with a higher number of unshrouding rivals. When unshrouding is costly, the relationship between the number of firms and unshrouding becomes non-monotonic. This is because stronger competition erodes unshrouding profits, incentivising firms to save on the fixed unshrouding costs when competition is too intense.

Wenzel (2015) studies how shrouding decisions are affected by differences in add-on production costs when, as in Wenzel (2014), the education effect of unshrouding depends on the number of unshrouding firms. The author shows that firms with low add-on production costs have a stronger incentive to unshroud, as they benefit more from an increase in add-on sales. For naïve consumer populations of intermediate size, the asymmetric unshrouding incentives permit a partial shrouding equilibrium, in which efficient firms unshroud and inefficient firms shroud.

In Murooka (2015), the decision whether to educate consumers about hidden add-on fees is taken by an intermediary under an agency model. The author shows that intermediaries have a weaker incentive to educate consumers the larger the hidden fee is, as the hidden fee finances higher commissions. As a result, intermediaries fail at their role to educate consumers and deception (that is, provision of products with hidden fees) emerges in equilibrium precisely when hidden fees are large.

Related to the literature on shrouding add-ons and additional fees are various works that study firms' incentives to hide product information or confuse consumers in order to soften competition on the visible prices. Some works study the use of different price frames (e.g. Piccione and Spiegler, 2012; Chioveanu and Zhou, 2013; Spiegler, 2014) or price complexity (Carlin 2009; Carlin and Manso, 2011) by firms to make product comparison difficult. Others (e.g. Ellison and Wolitzky, 2012; Wilson, 2010) study firm efforts to increase search costs in order to deter consumers from searching for better deals. Spiegler (2016) provides an excellent review of the literature on choice complexity and consumer confusion.

1.2.4 Externalities between naïve and sophisticated consumers

When consumers naïvely underestimate their demand for add-ons (or, equivalently, their probability of paying additional fees and charges), a relevant question is whether sophisticated consumers can protect them from being exploited. In some settings, sophisticated consumers can improve the deals available to all consumers by inducing firms to compete in all product dimensions. Armstrong (2015) refers to this competition-strengthening effect that sophisticated types create as a "search externality".

A search externality takes place in Armstrong and Chen (2009). In this model, firms can offer high quality (e.g. a low add-on fee) at high marginal cost and low quality (e.g. a high add-on fee) at zero cost. Sophisticated consumers observe both prices and qualities, while naïve consumers only observe prices, believing that quality is always high. In a market populated solely by naïve types, all firms supply low quality. Introducing some sophisticated types incentivises firms to randomise between high and low qualities. Interestingly, the relationship between consumer surplus and the proportion of sophisticated types is non-monotonic. This is because randomisation across qualities introduces price dispersion. The surplus loss associated with sometimes selling low qualities at high prices to naïve types dominates the surplus gain from the offering of high qualities when the fraction of sophisticates is relatively small.

A search externality can also take place in markets in which sophisticated types cannot avoid add-on fees, and upfront prices of all firms are competed down to a price floor. With upfront prices equal to the floor, naïve types do not exert any more competitive pressure and get shared evenly across firms. Sophisticated types, on the other hand, incentivise firms to compete on add-on fees. Interpreting his framework as one in which upfront prices are fixed and firms compete over add-on fees, Schultz (2004) provides insights on how naïvete about add-on fees affects horizontal differentiation in the market. He shows that product differentiation increases in the number of naïve consumers, as the weaker competitive pressure due to naïvete allows firms to translate product differentiation into even higher add-on fees.

The model of Chapter 4 investigates the relationship between naïvete and horizontal product differentiation when search externalities are present in a more general setting, revealing parameters for which Schultz's finding applies and results in intermediate product differentiation, and parameters for which other forces result in maximum, or minimum product differentiation.

If competition intensity is sufficiently intense, search externalities can result in price dispersion in a fashion similar to Varian (1980). Intuitively, if rivals' fees are low enough, firms have an incentive to set high add-on fees and exploit "captive" naïve types. The model of Chapter 4 features parameter values for which such price dispersion takes places, with unique results on product differentiation incentives.

Search externalities do not always arise. If sophisticated types avoid add-ons they do not intensify competition on add-on fees. Instead, they take advantage of the lower base good prices that are funded solely by naïve types' add-on purchases. Armstrong (2015) refers to such cross-subsidisation from naïve to sophisticated types as a "rip-off" externality. Armstrong and Vickers (2012) illustrates this externality type in a model of personal current accounts. Cross-subsidisation between consumers that are inattentive about past mobile phone usage and consumers that are attentive takes place in Grubb (2015b). The model in Chapter 3 of this thesis exhibits such rip-off externalities and shows how their emergence affects product differentiation incentives.

1.2.5 Distortions in markets with add-ons and naïve consumers

When naïve consumers underestimate future add-on payments, distortions can emerge. An avoidance distortion arises when consumers must spend effort to avoid high add-on fees. For instance, in Grubb (2015b) consumers need to spend attention costs in order to ensure that they do not incur overage charges. In Armstrong and Vickers (2012), avoiding overdraft charges requires a diligence cost. In Armstrong *et al.* (2009), consumers need to spend search costs to looks for firms that offer better deals. Avoidance costs are not always wasteful; Zenger (2013) shows that some avoidance effort may be efficient if add-on production costs are high.

A participation distortion appears when firms use add-on revenues to subsidise base good prices. This subsidisation results in consumers buying too many base goods. For example, in Zegners and Kretschmer (2017) if aftermarket power is sufficiently high then the base good subsidy is large enough that even consumers with valuation below marginal cost buy the base good. Heidhues and Koszegi (2015) demonstrates such price distortions in the US credit market. Crucially, the extent of this participation distortions depends on demand and supply elasticities, as mentioned in Grubb (2015a). Under certain circumstances, the participation distortion may be beneficial. For instance, using an isoelastic demand curve, de Meza and Reyniers (2012) show that the participation distortion can improve welfare by allowing consumers that value the base good to purchase it, when they would not do so in the absence of subsidies.

A particularly worrying distortion is that underestimation of future add-on payments can lead to the prevalence of low quality products. In Michel (2018), some consumers naïvely underestimate the costs of returning a product to a manufacturer when choosing whether or not to buy a warranty. If the size of warranty payments can signal product quality, consumer naïvete can permit a monopolist to offer a low-quality product with a warranty that signals high quality. Given sufficiently strong underestimation of return costs, low quality provision to naïve consumers can emerge even in a competitive market.

In the search model of Gamp and Krahmer (2018), “candid” firms, which produce a good of efficiently high quality at high cost, compete with “deceptive” firms, which produce a good of inefficiently low quality at low cost. Quality can be interpreted as presence or absence of hidden add-on fees or surcharges. Sophisticated consumers recognise low quality products, and search until they find a sufficiently cheap candid product. Naïve consumers, on the other hand, search for the cheapest product without regards to quality. In equilibrium, sophisticated consumers search and buy candid products, while naïve consumers buy the first product they see, and may end up with a deceptive product. As search frictions disappear, sophisticated consumers search between all candid firms, driving their mark-up to zero. Deceptive firms, on the other hand, maintain positive mark-ups due to their lower production costs. As a result, the deceptive business model dominates, and candid products are driven out of the market.

In Heidhues *et al.* (2017), as already discussed, shrouding attributes can permit firms profitably sell wasteful products. The authors extend their model to markets in which firms sell high-quality products at transparent prices, and low-quality products with add-on fees. In such markets, there exist equilibria in which firms sell the low-quality product with shrouded fees to naïve consumers, earning positive profits from it, and sell the high-quality product to sophisticated types, earning zero profits from it.

In a companion paper, Heidhues *et al.* (2016), the authors show that naïvete can also distort firms’ innovation incentives. When shrouding is an equilibrium, firms have no incentive to invest in *value-increasing* innovations if these innovations are *non-appropriable*, as rivals that copy the innovation may find it profitable to unshroud. However, the same is not the case with non-appropriable *exploitative* innovations – that is, innovations that increase the shrouded prices that naïve consumers are paying. A firm has an incentive to invest in such innovations because these innovations *(i)* allow the innovative firm to extract more surplus from naïve types, and *(ii)* allow all rivals that copy to extract more surplus as well, weakening their incentives to unshroud. Incentives for *appropriable* value-increasing innovations may also be distorted downwards under shrouding, as innovative firms consider the possibility that rivals unshroud in response to a successful innovation.

Chapters 3 and 4 reveal that naïvete with respect to add-on fees can introduce a novel type of distortion: a distortion in horizontal product differentiation. We show that distortion in this dimension may not be necessarily harmful; in fact, it may help correct the excessive product differentiation predicted by classical models (e.g. d'Aspremont *et al.*, 1979; Neven, 1985).

1.2.6 Can regulation protect naïve consumers?

When exploitation of naïve consumers is severe, regulators may intervene to improve market performance. Supply-side interventions, such as *price controls,* aim at directly constraining firms' exploitative behaviour. Demand-side interventions, such as *disclosure remedies,* aim at improving consumers' ability to choose good deals and avoid exploitation. Fletcher (2016) provides an excellent review and evaluation of various demand-side interventions in real markets.

Supply-side responses play an important role in the effectiveness of interventions. In markets with add-ons, interventions may succeed at reducing consumers' add-on expenses, but firms may adjust their business strategies and undo some of the benefits. Typically, firms may respond to a reduction in add-on revenues by increasing other prices. This is commonly known as a *waterbed effect*.

Grubb (2015b) shows the possibility for a waterbed effect in a market for calling plans. In his model, consumers that are not attentive with respect to past consumption make calls that exceed their plan's coverage and incur penalty charges. A bill-shock regulation that alerts consumers when they are about to exceed their plan coverage can protect naïve consumers that would otherwise incur penalty charges. The author shows that eliminating penalty expenses may not make the average consumer better off; in equilibrium firms respond by raising first-period prices.

The importance of price readjustments is stressed in an empirical work by Grubb and Osborne (2015), which simulates the effect of bill-shock alerts on phone plans using data from the years 2002-2004. They show that, holding other prices fixed, consumers benefit from the alerts. When equilibrium price readjustments are taken into account, the alerts may benefit or harm consumers, depending on whether consumers underestimate their probability on incurring penalty fees or not.

Agarwal *et al*. (2014) show that the extent of the waterbed effect depends on cost pass-through. From the firms' point of view, add-on revenues are equivalent to a reduction in marginal costs. As such, an intervention that limits add-on revenues is equivalent to a marginal cost tax, and its effect on base good prices depends on mechanisms well-studied in the pass-through literature (e.g. Weyl and Fabinger, 2013). An example case in which pass-through is zero takes place when base good prices are constrained by a price floor both before and after regulation takes place. Armstrong and Vickers (2012) show that this is a possibility in UK retail banking.

A difference-in-difference analysis by Agarwal *et al.* (2015) finds no support for a waterbed effect after the implementation of the CARD Act, a regulatory intervention aimed at protecting U.S. credit card users from incurring unanticipated charges. The authors find that the regulation reduced overall borrowing costs, and they no evidence of offsetting increases in interest charges or interchange fees, or offsetting reductions in credit volumes and various operational costs.

For a regulator with distributional concerns[9], a market intervention that limits add-on payments may be desirable even in the presence of a strong waterbed effect. If, for instance, it is only naïve consumers that incur add-on charges, eliminating these charges benefits naïve consumers at the cost of harming sophisticated consumers that only pay the base prices. The model in Chapter 3 of this thesis, describes a market situation in which interventions affect naïve and sophisticated consumers differently.

Firms' shrouding behaviour plays a role on whether regulatory interventions yield net benefits. By modifying the framework of Gabaix and Laibson (2006), Ko and Williams (2017) show that a price control remedy may harm both social surplus and the surplus of naïve consumers if it supports a shift from an unshrouding to a shrouding equilibrium. This is because a shrouding equilibrium features inefficient add-on substitution by sophisticated consumers, and high add-on expenses by naïve consumers. Using a similar model, Kosfeld and Schuwer (2017) demonstrate that disclosing hidden add-on fees can harm social welfare, unless the disclosure is strong enough to shift the market to an equilibrium in which add-ons are fully unshrouded. If

[9] Helping vulnerable consumers is at the core of the Competition and Markets Authority's 2018/2019 plan. See https://www.gov.uk/government/publications/competition-and-markets-authority-annual-plan-2018-to-2019/competition-and-markets-authority-annual-plan-20181.

the market features shrouding after the disclosure remedy takes place, then at least some of the newly educated naïve types prefer to (inefficiently) substitute away from add-on consumption.

Regulatory interventions may also affect consumers' own efforts to protect themselves from exploitation. In the search model of Armstrong *et al.* (2009), regulation affects consumer self-protection efforts in a detrimental way. Price controls dampen consumers' incentives to become informed, softening competition. By considering consumers' limited mental capacity, Heidhues *et al.* (2018) show that the effect of regulation on consumers' self-protection efforts is positive. In their model, consumers have to choose between studying a product's features in order to avoid hidden charges and browsing more products in order to find a better deal. Regulating hidden charges permits consumers to devote more effort into shopping around, strengthening competition and improving welfare.

The models in Chapters 3 and 4 of this thesis contribute to the literature that explores the effectiveness of regulatory interventions in markets with add-ons by studying whether certain types of interventions are desirable when horizontal product differentiation is a choice variable. The two models reveal a novel effect; regulatory interventions can alter firms' equilibrium choice of product differentiation. Whether this is to the benefit or detriment of consumer and social surplus depends on pre- and post-intervention market parameters.

Chapter 2

Add-on Pricing Equilibrium with Rational Consumers

2.1 Introduction

In a wide range of industries, firms charge extra for add-ons that complement their core products. Hotels sell breakfast or internet access separately; airlines charge extra for on-board meals or luggage; smartphone apps sell functionality or cosmetic upgrades. Add-ons are a major source of revenue for many of these firms. Apps following the freemium business model rely entirely on add-on revenues. Low-cost airlines are reported to earn up to 46% percent of their total revenue from ancillary services.[10] Despite the popularity of "add-on pricing" as a business strategy, identifying the incentives behind firms' choice to unbundle and offer base goods and add-ons separately remains an open question. This chapter attempts to shed some light on this question.

While the firms often enjoy monopoly power on the add-ons they sell, traditional competition arguments (Shapiro, 1995; Lal and Matutes, 1994; Verboven, 1999)

[10] https://www.telegraph.co.uk/travel/news/airlines-that-rely-most-on-extra-charges/ [accessed August 2018].

suggest that add-on revenues have no effect on profits; firms compete away add-on revenues by offering low (even below-cost) prices for the base goods. Using a competitive price discrimination model, Ellison (2005) shows that add-on pricing may increase profits if an adverse selection problem prevents firms from cutting base good prices below a certain threshold. However, he also shows that add-on pricing is not an equilibrium strategy precisely when such a threshold is binding.[11]

In this chapter we model two market settings in which unbundling base goods from add-ons (what we call the *add-on pricing strategy*) may be an equilibrium even when adverse selection constrains competition on base prices. In each market setting we specify conditions under which the adoption of an add-on pricing strategy is individually rational and permits firms to enjoy supra-competitive profits. In contrast with much of the extant literature, our work does not rely on behavioural assumptions regarding consumer choices; consumers correctly anticipate own behaviour and have rational expectations about unobservable features.

We begin in section 2.2 by presenting a benchmark duopoly model to simply illustrate Ellison's insight: that an add-on pricing equilibrium does not exist when a constraint prevents add-on pricing firms from cutting base prices below a certain threshold. Following existing literature (e.g. Heidhues *et al.*, 2017), we model this constraint as an exogenous price floor. In real markets, such a price floor can be established if low base good prices attract a disproportionate amount of low-value consumers that refrain from add-on purchases.[12] For example, setting flight ticket prices below a certain threshold may fill airplane seats with consumers that prefer to substitute away from on-board meals. Selling apps for negative prices may attract arbitrageurs that enjoy the subsidy without using the app itself. The price floor, when

[11] In section 2.2 we illustrate Ellison's mechanism in a simple benchmark model. In Ellison's original work, add-on pricing is defined as the business strategy in which firms sell add-ons at high unadvertised prices. An adverse selection problem arises because competition on base prices attracts a disproportionate amount of low-type consumers that do not buy the add-ons. While this dampens price-cutting incentives, a profitable deviation from an add-on pricing equilibrium takes place if the choice to advertise any price is endogenous. When this is the case, any firm can bypass the adverse selection problem by advertising a slightly lower add-on price and attracting high-type consumers.

[12] This is precisely the adverse selection problem that prevents price cutting in Ellison (2005). See Ellison and Ellison (2009) for evidence of an adverse selection problem in a market for computer parts.

binding, permits add-on pricing firms to earn supra-competitive profits. However, an equilibrium in which both firms follow the add-on pricing strategy does not exist in this case. Firms can deviate to offering an integrated good (what we call the *all-inclusive strategy*) and profitably undercut.

In section 2.3 we develop our first main model. In this model firms that follow the add-on pricing strategy can condition the production of an add-on on its purchase and save on production costs when an add-on purchase does not take place. This specification is particularly relevant to real markets with physical add-on products. For example, hotels can save on ingredients costs if residents do not buy breakfast; airlines can save on handling and fuel costs if passengers do not check luggage; supermarkets can save on transport costs if consumers visit the store instead of buying online.

We show that the prospect of abandoning potential cost savings associated with the add-on pricing strategy may be sufficient to deter firms from bundling base goods and add-ons (all-inclusive strategy) and undercutting, even when a binding price floor keeps base good prices relatively high. Compared to a world in which both firms bundle base goods with add-ons, the add-on pricing equilibrium, whenever it exists, improves social surplus and may or may not benefit consumers.

Our second main model, in section 2.4, analyses a market in which products may or may not match consumers' tastes. Informed consumers know the match quality of each product in advance and purchase from the cheapest firm that matches their tastes. Uninformed consumers are uncertain about match quality, and only learn about the match quality of a given product after purchase. Add-on pricing allows uninformed consumers to learn about match quality by buying only the base good, thus permitting them to condition purchase of the add-on on the revelation of good matching.

This model specification is applicable to various markets with experience goods, and in particular to digital products. For example, smartphone users that download a freemium app can taste some basic product features for free and may buy optional functionality or cosmetic upgrades after good matching is realised. Purchasing a paid app, instead, is riskier if a user is uninformed about match qualities.

We show that the voluntary nature of add-on purchases allows add-on pricing firms to price discriminate between informed and uninformed consumers. In contrast, all-inclusive firms can only offer the same deal to both consumer types. Under certain parameter specifications, the prospect of abandoning the ability to price discriminate can deter firms from following the all-inclusive strategy in order to undercut. When the price floor constraint is binding, this can establish an add-on pricing equilibrium with supra-competitive profits. Our model shows that this is the unique pure strategy equilibrium when matching is uncertain. This may, in part, explain the prevalence of the freemium business model in the apps market. Compared to a world in which both firms bundle base good with add-ons, the add-on pricing equilibrium, whenever it exists, improves social surplus and firm profits.

This chapter is, to the best of our knowledge, the first work that shows existence of an add-on pricing equilibrium even when a price floor limits add-on firms' ability to cut prices and even when consumers are fully rational. Our two main models are applicable to a wide range of industries in which add-on pricing is a popular business strategy.

Several related works study firms' incentives to unbundle base goods from add-ons, but do so without addressing the possibility that a price floor may constrain price cutting. Closest to our cost-savings model is Geng and Shulman (2015), who study how cost savings associated with add-on pricing affect competitive firms' choice of business strategy when consumers are heterogeneous in their price sensitivity. Lin (2017) shows that vertical product differentiation plays a role in the firms' choice to unbundle or integrate. Geng *et al.* (2018) show that an upstream monopolist's choice of business strategy depends on the distribution contract this firm has with a downstream platform. Cui *et al.* (2018) study how the monopolist's incentives to unbundle ancillary services changes depending on whether that monopolist can price discriminate between high and low-value consumers on the main service.

In a related question, Gomes and Tirole (2018) explain why some markets feature high add-on prices, while in other markets add-ons are given away for free. Modelling a market in which consumers may choose not to buy the base good if add-on prices are too high, they show that the possibility to miss base good sales deters a base good monopolist from setting high add-on prices and may even incentivise the monopolist to offer add-ons below cost. The higher the mark-up on the base good, the weaker the incentive to set high add-on prices. Hold-ups can occur when consumers are uninformed about add-on prices or naïvely ignore them; however, the extent of the hold-up is limited by the concern for missed sales. Addressing the same question in an earlier work, Fruchter *et al*. (2011) show that a monopolist may have an incentive to sell add-ons at a loss to low-value consumers, as the associated additional surplus that these consumers enjoy may allow the firm to charge higher base prices to both low and high-value consumers.

Add-on pricing is, at its core, a second-degree price discrimination strategy, aimed at extracting higher surplus from high-value consumers. In markets with rational consumers, the price discrimination is preference-based: add-ons are bought by consumers that value them the most. Stole (2007) provides an excellent review of classical preference-based price discrimination models of imperfectly competitive markets.

In markets with boundedly rational consumers, add-on pricing can permit a second-degree naïvete-based price discrimination.[13] This type of price discrimination is studied in a large literature following Gabaix and Laibson (2006). In that literature, firms can choose between shrouding add-ons to extract large surplus from naïve consumers that accidentally purchase them, or unshrouding and selling add-ons to both naïve and sophisticated consumers at lower prices. With regards to its ability to discriminate, shrouding is similar to unbundling and selling add-ons separately, while unshrouding is similar to selling add-ons to everyone at an all-inclusive price. In the absence of a binding price floor, firms' incentive to keep add-ons shrouded depends, among other parameters, on the proportion of behavioural consumers in the market (Gabaix and Laibson, 2006), on competition intensity (Wenzel, 2014), and on cost asymmetries (Dahremoller, 2013; Wenzel, 2015). With a binding price floor, however, shrouding is an equilibrium only if it permits firms to sell to naïve consumers at prices above the monopoly value (Heidhues *et al.* 2017).

The add-on pricing strategy is related to drip pricing, a business strategy in which firms reveal various charges at the later stages of a transaction. Shelanski *et al.* (2012) provides a review of this business strategy from a consumer policy perspective. When

[13] For a study of first and third-degree naïvete-based price discrimination, see Heidhues and Koszegi (2017).

the dripped price is optional (e.g. purchase of one-day delivery), drip pricing closely resembles add-on pricing. When the dripped price is mandatory, the business strategy is also referred to as partitioned pricing. Partitioned pricing differs from the standard process of selling by posting a single price due to psychological biases that often arise from the revelation of costs later in the purchasing process (Ahmetoglu *et al.*, 2014; Office of Fair Trading, 2013; Greenleaf *et al.*, 2016).

2.2 Benchmark model

In this section we develop a benchmark model to simply illustrate the finding of Ellison (2005) that an add-on pricing equilibrium does not exist if the add-on pricing strategy constrains firms' ability to offer low prices.

2.2.1 Model structure

Consider two identical firms selling a good of quality v to a consumer segment of size 1. Consumers are homogeneous and have unit demand. Firms face marginal production cost c and compete on prices. Each firm i operates under a business strategy $R_i \in \{I, A\}$. Strategy I, or *all-inclusive* strategy, involves selling the good in one piece, at an observable *all-inclusive price* p_I. Strategy A, or *add-on pricing* strategy, involves separating the product in two parts; a base good of quality $v - w > 0$ at an advertised *base price* p_b and an add-on component of quality w at an unadvertised[14] *add-on price* p_a. p_b is assumed to be bound below by a price floor f, where $f \leq c$. Consumers can purchase the add-on component only from the firm from which they purchased the base good and only after they purchase the base good. They

[14] Advertising add-on prices may be prohibitively expensive. For example, advertising optional extras available with a car purchase takes up valuable space in a specialist magazine, or valuable time in a TV commercial. Advertising an add-on price may also be ineffective at informing consumers about the associated add-on expense. For example, informing a consumer of a videogame about the price of an in-app purchase that gives her an extra life does not convey any information about how often she will need that extra life.

can observe the p_a of their firm of choice only after they purchase the base good from that firm. They have rational expectations regarding prices they cannot observe.

Consumers that purchase the base good from a firm also purchase the add-on component from that firm at any price $p_a \leq w$ with probability $q \in (0,1)$. [15] A firm that follows the add-on pricing strategy has no incentive to cut add-on prices since they are unadvertised, and thus always sets $p_a = w$. Consumers and firms know the value of q but do not know ex ante whether a given consumer will buy the add-on. For brevity, we refer to firms following the add-on pricing strategy as *add-on firms*, and firms following the all-inclusive strategy as *all-inclusive firms*.

Consumers purchase from the firm that offers the best deal. To resolve ties, we assume that firms share consumers evenly if they offer the same deal. We also assume that consumers always buy if they are indifferent between buying and not buying.

Timing

Firm decisions take place in two stages. In stage 1 firms simultaneously choose business strategies. In stage 2 they simultaneously choose prices. A firm that chooses business strategy I sets and posts an all-inclusive price, p_I. A firm that chooses business strategy A sets a price pair (p_b, p_a) and posts its base price, p_b. Consumers observe the pricing strategies of each firm and the posted prices, form rational expectations about the price of the add-on component, if any, and choose from which firm to purchase. We solve the game using backward induction and look for symmetric subgame perfect pure strategy equilibria. We refer to a pure strategy equilibrium in which both firms follow the add-on pricing (all-inclusive) strategy as an *add-on pricing (all-inclusive) equilibrium*.

[15] This assumption can be interpreted in the following way. Each consumer k has individual valuation for the add-on component w_k, where $w_k = w$ with probability q and $w_k = 0$ with probability $1 - q$.

2.2.2 Equilibrium analysis - Non-existence of an add-on pricing equilibrium

We first consider the subgame in which both firms follow the add-on pricing strategy (*add-on subgame*). In this subgame, firms earn, on average, wq from each customer and compete on base prices.

If the price floor is not binding then the equilibrium base price is such that profits are equal to zero; that is, $p_b^* = c - wq$. The intuition behind why firms set below-cost prices in equilibrium is that firms have an incentive to compete more fiercely for consumers because they expect revenues wq from add-on purchases. This more intense competition results in firms subsidising base prices enough to dissipate all add-on revenues, leaving firms with zero profits. Due to this subsidisation on base prices, consumers pay an expected total price $E_{nf}^{add} = p_b^* + wq = c$, where subscript nf indicates a non-binding price floor, and the superscript *add* indicates the add-on subgame.

If the price floor is binding (that is, if $f > c - wq$) firms cannot compete away all add-on revenues by subsidising lower base prices, as the price floor limits price cutting. As a result, equilibrium base price is $p_b^* = f$, and expected total price is $E_f^{add} = f + wq$ (where subscript f denotes a binding price floor). Benefitting from the price floor constraint, firms share consumers evenly and earn supra-competitive profits: $\Pi^{add} = \frac{f+wq-c}{2}$. Note that base prices are partially subsidised, as firms' incentive to compete more fiercely for the add-on revenues remains. Note also that the subsidisation is precisely why the price floor constraint binding; assumption $f \leq c$ implies that the price floor can be binding only if firms set below-cost prices.

To see that a symmetric add-on pricing equilibrium does not exist when the price floor is binding, consider the subgame in which firm i follows the add-on pricing strategy and firm j follows the all-inclusive strategy (*asymmetric subgame*). In this case, the all-inclusive firm can always undercut its add-on pricing rival. To see why, notice that an all-inclusive firm's pricing behaviour is not constrained by the price floor. This is because it expects no future add-on revenues, so it has no incentive to offer a subsidised price. Thus, its pricing behaviour is only constrained by the marginal cost (c). If the price floor is binding, then c is greater than the lowest expected total price that an add-on firm can offer ($f + wq$). Thus, the all-inclusive firm can always offer a better deal to consumers. In the unique pure strategy price equilibrium of the asymmetric subgame, the add-on firm sets a base price $p_b = f$, with an associated expected total price $E_f^{asym} = f + wq$ (superscript *asym* indicates the asymmetric subgame, and f indicates a binding price floor). The all-inclusive firm sets an all-inclusive price p_I slightly below $f + wq$ and serves every consumer, earning $\Pi_I^{asym} = f + wq - c$.

Since $\Pi_I^{asym} > \Pi^{add}$, each firm is better off choosing the all-inclusive strategy when its rival chooses the add-on pricing strategy. Hence, a symmetric pure strategy equilibrium in which both firms follow the add-on pricing strategy does not exist.

On the contrary, an all-inclusive pure strategy equilibrium always exists. In the subgame in which both firms follow the all-inclusive strategy (*all-inclusive subgame*) firms compete a la Bertrand, setting $p_I = c$ and earning $\Pi^{inc} = 0$, where superscript *inc* indicates the all-inclusive subgame. No firm has incentive to deviate to the add-on pricing strategy, as the firm that follows the add-on pricing strategy in the asymmetric subgame also earns zero profits.[16]

An add-on pricing equilibrium trivially exists when the price floor is non-binding. In the asymmetric subgame, a non-binding price floor permits the add-on firm to subsidise the base price enough to match any deal that the all-inclusive firm may offer. As a result, in the equilibrium of this subgame, firms share consumers and earn zero profits each. The add-on firm dissipates future add-on revenues by setting $p_b = c - wq$, resulting in an expected total price $E_{AI}^{nf} = c$. The all-inclusive firm sets $p_I = c$. Thus, an add-on pricing equilibrium exists when the price floor is non-binding since a firm always earns, regardless of its own business strategy, zero profits when its rival chooses the add-on pricing strategy. Similarly, an all-inclusive equilibrium in which both firms earn zero profits also exists. Proposition 1 summarises the equilibrium outcome in the benchmark model.

[16] Firms have no incentive to deviate if a deviation is not strictly profitable.

Proposition 1. (Ellison benchmark).

(i) In the benchmark model, if the price floor constraint is non-binding ($c \geq f + wq$)

(a) There always exists an add-on pricing equilibrium.

(b) There always exists an all-inclusive equilibrium.

(ii) In the benchmark model, if the price floor constraint is binding ($c < f + wq$) there always exists an all-inclusive equilibrium. This is the unique symmetric equilibrium in pure strategies.

Proof: *In appendix A.*

Note that the results and intuitions do not change if pricing strategies and individual prices are chosen at a single stage (as in Ellison (2005)). With a binding price floor, a firm is always better off selling through an all-inclusive strategy if its rival follows the add-on pricing strategy, as it can bypass the price floor constraint and undercut. This only permits an all-inclusive equilibrium. With a non-binding price floor, an add-on firm can match any price offered by an all-inclusive firm, permitting an add-on pricing equilibrium with zero profits. Similarly, an all-inclusive equilibrium with zero profits also exists.

This simple benchmark model illustrates the intuition in Ellison (2005) that a symmetric pure strategy equilibrium in which both firms follow the add-on pricing strategy does not exist when that business strategy constrains the adopting firm's ability to compete by offering lower total prices.[17] Firms can profitably deviate to the all-inclusive strategy, bypass the price floor constraint and offer a strictly better deal to consumers. In the following sections we modify our benchmark model to settings that more closely describe real markets in which add-on pricing strategies are prevalent. We show analytically that the firms' incentive to deviate is weaker in these settings and describe conditions under which an add-on pricing equilibrium arises.

[17] Our model is a simple variant, developed solely to capture this intuition. The model in Ellison (2005) is much richer.

2.3 The cost-savings model

In this section we modify the benchmark model in order to study markets in which add-on pricing firms can condition the production of the add-on component on its purchase. This can permit firms to save on costs when a consumer has no demand for the add-on. Our cost-savings model specification applies to a range of markets in which firms often offer add-on services. Hotels charge separately for breakfast and save on production costs if residents do not purchase it; airlines charge extra for checked luggage and can save on handling and fuel costs if passengers only travel with carry-on luggage; supermarkets charge more for home deliveries, saving on transport costs when consumers visit the store for groceries.

We show that the associated cost savings that a firm can enjoy by following an add-on pricing strategy can mitigate firms' deviation incentives. Under certain conditions, this may establish a symmetric add-on pricing equilibrium even when that equilibrium is characterised by positive firm profits.

2.3.1 Model structure

We modify the benchmark model as follows. We assume that a firm that follows the add-on pricing strategy incurs marginal cost c_b to produce and sell the base good, and a separate marginal cost c_a to produce the add-on component. The production of the add-on component only takes place if a consumer buys the add-on. To ensure that an add-on firm always has an incentive to offer the add-on component, we assume that $w > c_a$. Following the benchmark model, add-on prices are unadvertised, so firms set $p_a = w$ in any equilibrium in which they follow the add-on pricing strategy. Also following the benchmark model, consumers buy the add-on only with probability q. Hence, an add-on firm's marginal cost function is given by $C_A = c_b + c_a q$. As in the benchmark model, a firm that follows the all-inclusive strategy faces a constant marginal cost c.

In order to avoid trivial cases in which one business strategy is always less costly than the other, we assume $c_b + c_a \geq c \geq c_b$;[18] that is, a firm that sells both price components to a consumer would save on costs by following the all-inclusive strategy. However, if the consumer only has demand for the base good then the seller can save on costs by following the add-on pricing strategy. As in the benchmark model, we solve the game using backward induction, and look for symmetric subgame perfect pure strategy equilibria. All other model assumptions are as in the benchmark model.

2.3.2 Equilibrium analysis

With different marginal cost structures, the firms' choice of business strategy is governed by two incentives. The first is the incentive to offer the best deal to consumers. Since products are homogeneous, the firm offering the best deal serves the whole market. The second is the incentive to take advantage of cost savings in order to enjoy the highest possible profit for any given price.

In absence of a binding price floor constraint, both incentives work in the same direction. The business strategy that features the lowest costs also permits firms to offer the best deal to consumers, and is, thus, chosen in equilibrium. With a binding price floor this is not necessarily the case. For certain parameter values, firms offer the best deal to consumers when following the all-inclusive strategy but enjoy the lowest costs when following the add-on pricing strategy. This may permit an add-on pricing equilibrium to emerge even if that equilibrium features worse deals for consumers. Proposition 2 presents the symmetric pure strategy equilibria that arise from the interaction of the two incentives.

[18] The inequality in the left side of the expression may hold if selling add-on separately results in higher operating costs. For example, selling in-flight meals instead of giving them away requires marketing efforts to induce flyers to purchase. Nevertheless, the inequality is not necessary for our key results. On the contrary, our key results arise more easily if we restrict our analysis to cases with $c_b + c_a = c$.

Proposition 2. (Pure strategy equilibrium in the cost-savings model).

(i) *In the cost-savings model, if the price floor constraint is non-binding ($c_b + c_a q \geq f + wq$)*

(a) There exists an add-on pricing equilibrium if and only if $c - c_b - c_a q \geq 0$

(b) There exists an all-inclusive equilibrium if and only if $c - c_b - c_a q \leq 0$

(ii) *In the cost-savings model, if the price floor constraint is binding ($c_b + c_a q < f + wq$):*

(a) There exists an add-on pricing equilibrium if and only if $c - c_b - c_a q \geq f + wq - c$.

(b) There exists an all-inclusive equilibrium if and only if $f + wq - c > 0$

Proof: *In appendix A.*

Before addressing the logic behind the outcome described in Proposition 2, it is useful to clarify the conditions under which we consider the price floor to be binding in this model. We define the price floor as binding if it does not permit add-on pricing firms to compete away add-on revenues that they anticipate at the add-on pricing subgame.[19] Given a non-binding price floor, firms at the add-on pricing subgame compete on base prices till the total expected price is equal to marginal cost; that is they set $p_b^* = c_b + c_a q - wq$. A binding price floor prevents firms from offering such low base prices, satisfying $f > p_b^*$. Rearranging gives that the price floor is binding if $c_b + c_a q < f + wq$ and non-binding otherwise.

Proposition 2 shows that, given a non-binding price floor constraint, each individual firm chooses the business strategy that is associated with the lowest cost in order to offer the best possible deal to consumers. As such, the conditions for which each equilibrium arises are mutually exclusive. $c - c_b - c_a q > 0$ indicates that the add-on pricing strategy is associated with lower total cost and thus offers the best deal

[19] A price floor that constrains add-on firms' behaviour at the add-on pricing subgame may not constrain the add-on firm's behaviour at the asymmetric subgame.

to consumers. When this is the case, there exists a unique add-on pricing equilibrium. Firms earn, on average, wq from each consumer, and set base prices low enough that they dissipate add-on revenues; that is, they set $p_b^* = c_b + c_a q - wq$. Equilibrium profits are zero. $c - c_b - c_a q < 0$ indicates that it is the all-inclusive strategy is associated with lower total cost and thus offers the best deal to consumers. In this case both firms choose that business strategy and compete a la Bertrand, setting $p_I^* = c$, sharing consumers evenly and earning zero profits each. In the knife-edge case in which both strategies have the same cost ($c - c_b - c_a q = 0$) both equilibria exist, as both strategies can offer the same deals to consumers.

The conditions determining the market equilibrium differ crucially when the price floor constraint is binding, as the binding price floor breaks the relationship between deals available to consumers and costs. Together with the fact that the two strategies face different marginal costs, this means that the market equilibrium is determined not only by which business strategy can offer the best deal to consumers, but also by the costs that each firm incurs to serve the consumers it attracts. To give an example, the best deal that add-on pricing firm can offer to consumers is an expected total price $f + wq$, while facing a marginal cost $c_b + c_a q$. This implies that an all-inclusive firm with marginal cost $c < f + wq$ can offer a better deal to consumers, but face higher marginal costs if, at the same time, $c_b + c_a q < c$. This results in equilibrium conditions that are overlapping for certain parameter values – for certain parameter regions both an all-inclusive and an add-on pricing equilibrium exist.

To see the equilibrium outcome analytically, recall that the best deal an all-inclusive firm can set is c, while the best deal an add-on firm can set is an expected total price $f + wq$. In the add-on subgame both firms compete on base prices, charging $p_b^* = f$ in equilibrium and sharing consumers. Equilibrium profits are $\Pi^{add} = \frac{f + wq - c_b - c_a q}{2}$. In the all-inclusive subgame, firms compete on the all-inclusive price, charging $p_I^* = c$ and earning $\Pi^{inc} = 0$.

In the asymmetric subgame, the add-on firm can serve all consumers by offering a better deal if $c > f + wq$. This results in profits $\Pi_A^{asym} = c - c_b - c_a q$ for the add-on firm, as it serves all consumers at a total expected price slightly below c, and $\Pi_I^{asym} = 0$ for the all-inclusive firm. Since $\Pi^{add} > \Pi_I^{asym}$, each firm is better off choosing the add-on pricing strategy when its rival chooses the add-on pricing strategy. Thus, there exists an add-on pricing equilibrium if $c > f + wq$. This inequality is satisfied in the condition of part *(ii-a)* of Proposition 2.[20]

The all-inclusive firm can serve all consumers by offering a better deal at the asymmetric subgame if $c < f + wq$. This results in profits $\Pi_I^{asym} = f + wq - c$ for the all-inclusive firm, as it serves all consumers at a price slightly below $f + wq$. The profits of the add-on firms are $\Pi_A^{asym} = 0$. Since $\Pi_A^{asym} = \Pi^{inc}$, a firm is indifferent between the two strategies if its rival follows the all-inclusive strategy. Thus, there exists an all-inclusive equilibrium if $c < f + wq$, as described in part *(ii-b)* of Proposition 2. However, this may not be the unique symmetric equilibrium. Comparison between Π_I^{asym} and Π^{add} yields that $\Pi^{add} > \Pi_I^{asym}$ if $c - c_b - c_a q > f + wq - c$, meaning that an add-on pricing equilibrium exists when this inequality is true. This is exactly the condition in part *(ii-a)* of Proposition 2.

Intuitively, the fact that an all-inclusive strategy may be able to offer a better deal to consumers is not sufficient to guarantee that an add-on pricing equilibrium does not exist. The add-on pricing equilibrium may still exist if the lower costs that a firm abandons by deviating to the all-inclusive strategy (that is, the difference between c and $c_b + c_a q$) can more than offset the higher revenues that this firm can earn by serving all consumers.

The coexistence of equilibria for certain parameter values raises the issue of equilibrium selection. Naturally, firms prefer to coordinate towards the add-on pricing equilibrium, as it yields positive profits ($\Pi^{add} = \frac{f + wq - c_b - c_a q}{2}$), unlike the all-inclusive equilibrium which yields zero profits.

As in the benchmark model, the equilibrium outcome does not change if pricing strategies and individual prices are chosen at a single stage. With a non-binding price floor, whether each equilibrium exists depends solely on which business strategy can offer the best deal to consumers. With a binding price floor, existence of an add-on pricing equilibrium depends not only on whether a rival firm can undercut by deviating to the all-inclusive strategy, but also on whether it abandons significant cost savings

[20] If $c > f + wq$ then the term $c - c_b - c_a q$ is positive when the price floor is binding.

in order to deviate to that strategy. Existence of an all-inclusive equilibrium depends solely on whether an add-on pricing firm can undercut its all-inclusive rival.

2.3.2.1 Comparative statics

Given a non-binding price floor, the equilibrium is determined solely by which business strategy features the lowest marginal costs. As such, higher marginal costs at the add-on pricing strategy (higher $c_b + c_a q$) and lower marginal costs at the all-inclusive strategy (lower c) facilitate the existence of an all-inclusive equilibrium and reduce the parameter space for which an add-on pricing equilibrium exists.

Given a binding price floor, higher per customer revenues at the add-on pricing strategy $(f + wq)$ decrease the parameter space for which an add-on pricing equilibrium exists and increase the parameter space for which the all-inclusive equilibrium exists. Intuitively, this is because higher $f + wq$ increase the profits that any firms can earn by deviating to the all-inclusive strategy. Higher $c_b + c_a q$, and lower c work in the same direction as $f + wq$, increasing the incentives of firms to deviate to the all-inclusive strategy.

The role of the add-on component's demand (q) is critical in determining the market equilibrium. Given a non-binding price floor, a higher q affects the equilibrium outcome only by increasing the marginal cost that an add-on firm faces from the production of the add-on component (higher $c_a q$). Hence, higher q reduces the parameter space for which an add-on pricing equilibrium exists and increases the parameter space for which an all-inclusive equilibrium exists.

For a market with a binding price floor, increasing q has two effects. First, it increases the add-on revenues that an add-on firm can earn from each customer it serves (higher wq). Second, it increases add-on firms' marginal costs (higher $c_a q$). Since $w > c_a$, increasing q has a positive net effect on the add-on firms' profits at an add-on pricing equilibrium is positive (to see this notice that Π^{add} is increasing in q). However, higher q also reduces the parameter space for which the add-on pricing equilibrium exists, as it increases the profits that a firm can enjoy by deviating to the all-inclusive strategy. This happens for two reasons. First, a higher q increases total expected prices at the add-on pricing equilibrium (higher $f + wq$), as it increases the likelihood in which they purchase the add-on component. This permits the deviating firm to set a higher p_I at the asymmetric subgame, thus increasing the profitability of a deviation to the all-inclusive strategy. Second, by increasing marginal costs at the add-on pricing strategy, a higher q reduces the cost-saving that a firm abandons in order to deviate to the all-inclusive strategy.

2.3.2.2 Welfare

In a market with a non-binding price floor, the business strategy used in equilibrium is the one that maximises social and consumer surplus. This is intuitive, as firms play the business strategy that minimises costs in order to offer the best possible deal to consumers. The relationship between social and consumer surplus and market equilibrium is not as straightforward in markets with a binding price floor.

The add-on pricing equilibrium, if it exists, always maximises social surplus. This is because the add-on pricing strategy requires strictly lower marginal costs than the all-inclusive strategy in order for firms to have no incentive to deviate. The all-inclusive strategy, on the other hand, may characterise an equilibrium even if that business strategy features higher marginal costs than the add-on pricing strategy. To see this analytically, notice that the all-inclusive strategy features lower marginal costs if $c < c_b + c_a q$, but an all-inclusive equilibrium exists if $c < f + wq$. Since, by the binding price floor condition we have $f + wq > c_b + c_a q$, we can see that an all-inclusive equilibrium exists despite being wasteful in terms of social welfare if $c_b + c_a q < c < f + wq$.

The all-inclusive equilibrium, whenever it exists, maximises consumer surplus. The intuition is straightforward: an all-inclusive equilibrium can only exist if the all-inclusive strategy features the lowest possible prices. On the contrary, existence of an add-on pricing equilibrium does not imply that the add-on pricing strategy features the best deal to consumers. This is because an add-on pricing equilibrium may exist even if the all-inclusive strategy permits lower prices, as long as the cost savings associated with the add-on pricing strategy are large enough to mitigate deviation incentives.

2.4 The matching uncertainty model

In this section we modify the benchmark model in order to study markets in which some consumers are uncertain about whether a certain product's characteristics matches their taste. Consumers that purchase from an add-on pricing firm can learn about whether they are a good match by buying the base good, and purchase the add-on only on the condition that good matching is revealed.

This model specification is applicable to various markets with experience goods, and in particular to digital products. Examples include smartphone apps and software-as-a-service (SaaS) products (e.g. Skype, Dropbox). Products in these categories often permit users to access some features for free, while charging them extra for functionality or cosmetic upgrades (such as the option to remove ads, improve the look of ones' in-game avatar or permit calls to landlines). Consumers can get a taste of the product through the basic version, and can buy the upgrades if the basic version suits their needs.

We show that the presence of consumers that are uninformed and uncertain about matching quality alongside consumers that are informed ex ante about matching quality can establish a profitable add-on pricing equilibrium by mitigating firms' incentives to deviate to the all-inclusive strategy and undercut. We also show that this is the unique symmetric pure strategy equilibrium that may arise in such a market; a pure strategy equilibrium in which both firms follow the all-inclusive strategy does not exist.

2.4.1 Model structure

As in the benchmark model, each firm i operates under a business strategy $R_i \in \{I, A\}$. Strategy I involves selling a good of quality v in one piece, at an observable all-inclusive price p_I. Strategy A involves separating the product in two parts; a base good of quality $v - w > 0$ at an advertised base price p_b and an add-on component of quality w at an unadvertised add-on price p_a. p_b is assumed to be bound below by a price floor f. We assume that both pricing strategies entail the same marginal cost, c.[21] To simplify modelling further, we also assume $c = f = 0$.[22]

Consumers can purchase the add-on component only from the firm from which they purchased the base good and only after they purchase the base good. They can observe the p_a of their firm of choice only after they purchase the base good from that firm. They have rational expectations regarding prices they cannot observe. Since add-on prices are not advertised, add-on firms have no incentive to compete on add-on prices, so they always set $p_a = w$.

The intrinsic utility a consumer enjoys from a specific product depends on how closely that product matches her taste. A product perfectly matches a consumer's taste with probability $q \in (0,1)$. If that is case the consumer enjoys the full utility that the product offers. If the product is a mismatch (probability $1 - q$), the product yields zero utility to that consumer. We refer to q as match quality. A consumer that buys from an add-on pricing firm buys the add-on component only if matching is correct.

All consumers know the value of q. However, they differ on whether they know in advance whether a specific product matches their taste. A proportion $s \in (0,1)$ of consumers are *informed* ex ante about match qualities. An informed consumer buys from the cheapest firm if both firms' products match her tastes (probability q^2). If only one firm's product matches her tastes (probability $q(1 - q)$) then she buys from that firm as long as she anticipates a total price below v. If no product matches her tastes (probability $(1 - q)^2$) then she does not make a purchase. Since she is informed about

[21] We do this for two reasons. First, assuming identical marginal costs avoids confounding any equilibrium effects arising from matching uncertainty with the effects of cost asymmetries that we studied in section 2.3. Second, the identical marginal cost assumption fits our primary example, digital experience goods, particularly well. Digital products typically have high fixed costs during the development stage, but almost zero costs for selling and distributing a copy to a buyer. It is, thus, unlikely that cost savings matter in a software firm's choice between all-inclusive or add-on pricing.

[22] Assuming a price floor equal to marginal cost is reasonable when marginal costs are zero, as, for example, in the case of digital goods. Negative prices may be undesirable because they may attract a large proportion of consumers that purchase the product solely to enjoy the subsidy and ignore any add-ons.

match quality before purchase, she always buys the add-on component if she purchases the base good.

A proportion $1-s$ of consumers are *uninformed* about match qualities ex ante. An uninformed consumer that buys from an all-inclusive firm learns about the match quality only after she purchases the product. As such, her ex ante intrinsic valuation of the all-inclusive product is vq. An uninformed consumer that buys from an add-on firm learns about the match quality after purchasing the base good. As such, her ex ante intrinsic valuation of the base good is $(v-w)q$. She buys the add-on component only if the base good purchase reveals that the product is a good match (probability q). If the product is a good match, then she also buys the add-on component at any $p_a \leq w$. Since she buys the add-on component with probability q, her intrinsic valuation of the complete product ex ante is vq. Since an uninformed consumer is unaware of match qualities in advance, she always buys from the cheapest firm.

We assume that firms choose business strategies and prices at a single stage. [23] A firm that chooses strategy I sets and posts an all-inclusive price, p_I. A firm that chooses strategy A sets a price pair (p_b, p_a) and posts its base price, p_b.

After business strategies and prices are chosen, consumers observe the business strategies of each firm and all posted prices, form rational expectations about the price of the add-on components, if any, and choose from which firm to purchase. Consumers who purchased the base good from an add-on firm choose whether or not to buy the add-on component.

As in the benchmark model, we look for symmetric pure strategy equilibria. We refer to a pure strategy equilibrium in which both firms follow the add-on pricing (all-inclusive) strategy as an *add-on pricing (all-inclusive) equilibrium*.

[23] Unlike in the models of the previous sections, we do not discuss the equilibrium of a 2-stage game in which firms first choose business strategies and then choose prices. This is because an equilibrium in which the two firms adopt different business strategies may not exist even in mixed strategies. While restrictive, the assumption of a setting is realistic in markets in which firms can costlessly switch between the two pricing strategies and notify consumers before purchase whether additional fees and charges apply.

2.4.2 Equilibrium analysis

Differences in the match quality information that is available to consumers can diminish the firms' price cutting incentive in two crucial ways. First, the presence of consumers that are matched with only one firm can make it profitable to set monopolistic prices and extract all the value from these consumers. In what follows we refer to these consumers as "captive consumers". We refer to consumers that pick from the cheapest firm as "contestable consumers".

Second, due to uncertainty in match quality, uninformed consumers' ex ante valuation of any product is less than the valuation of informed (and well-matched) consumers. This implies that in certain situations, serving uninformed consumers can require discrete price cuts, which may not be in any firm's best interest to provide.

Due to these two effects, firms' equilibrium behaviour goes beyond simply offering the best possible deal. Proposition 3 presents our findings regarding the existence of a pure strategy equilibrium.

Proposition 3. (Pure strategy equilibrium in the matching uncertainty model).

(i) *Let* $vsq(1-q) > w\left(sq(1-q) + \frac{sq^2}{2} + \frac{(1-s)q}{2}\right)$, *so that the price floor constraint is not binding in the matching uncertainty model. Then a symmetric pure strategy price equilibrium does not exist.*

(ii) *Let* $vsq(1-q) \leq w\left(sq(1-q) + \frac{sq^2}{2} + \frac{(1-s)q}{2}\right)$, *so that the price floor constraint is binding in the matching uncertainty model. Then there exists an add-on pricing equilibrium if* $\frac{1}{1+q} > s > \frac{1}{3(1-q)}$. *This is the unique symmetric equilibrium in pure strategies.*

Proof: *In appendix A.*

Proposition 3 shows that an equilibrium in pure strategies does not exist in absence of a binding price floor. The reason is the presence of captive consumers, and the logic of the result is as follows. As usual, prices which earn positive mark-ups cannot constitute a pure strategy equilibrium, as rivals can always profitably undercut.

However, prices that earn zero mark-ups cannot constitute a pure strategy equilibrium either. Regardless of whether firm i chooses all-inclusive or add-on pricing, given a sufficiently low total expected price by firm i, firm j's best response is not to undercut or match its rival's price, but to set a total expected price such that captive consumers are indifferent between buying and not buying. That is, its best response is to set either p_b such that $p_b + w = v$ (following the add-on pricing strategy), or $p_I = v$ (following the all-inclusive strategy), and serve only its captive consumers, earning $vsq(1-q)$. Since firms can always profitably deviate from any price pair, the non-binding price floor game exhibits price dispersion as in Varian (1980).[24]

This price dispersion requires us to redefine the condition under which we consider the price floor constraint to be binding. It is meaningless to define a binding price floor condition simply as a relationship which, when satisfied, prevents add-on firms from competing down to zero profits (as we did in the cost-savings model). This is because a lower bound that prevents such price cutting arises endogenously, as the lowest price below which firms find it profitable to price as monopolists and exploit their captive consumers. Therefore, a meaningful binding price floor condition needs to prevent add-on firms from cutting prices below this endogenous lower bound. In particular, that condition needs to ensure that no firm has incentive to deviate to a monopolistic price if both firms choose the add-on pricing strategy and set base prices equal to the price floor. Put formally, we define a price floor $f = 0$ as binding if it satisfies

$$w\left(sq(1-q) + \frac{sq^2}{2} + \frac{(1-s)q}{2}\right) \geq vsq(1-q) \qquad (BF)$$

(BF) is the binding price floor condition used in part *(b)* of Proposition 3 (and its reverse in part *(a)*). The left-hand side represents the profits of an add-on pricing firm that sets p_b equal to the price floor when its rival does the same. In this case, that firm serves all its captive consumers and half the contestable consumers in the market, earning w from each customer. The term $sq(1-q)$ represents the firms' captive

[24] Fully characterising a mixed-strategy equilibrium is challenging, as a firm can respond to a rival's action not only by choosing a specific total expected price but also by choosing a different business strategy. The characterisation of such an equilibrium lies outside the scope this chapter.

informed consumers, that is, the informed consumers that are only matched with that firm. The term sq^2 represents the informed consumers that are matched with both firms. Of these, the firm serves half. Finally, the term $(1-s)$ represents the uninformed consumers in the market. The firm serves half of these consumers. The proportion q of them buy the add-on, yielding profit w to the firm. The right-hand side represents the profits that any firm can earn by exploiting its captive informed consumers, by setting either $p_b + w = v$ (if following the add-on pricing strategy), or $p_I = v$ (if following the all-inclusive strategy).

Higher values of w and lower values of v increase the parameter space for which the binding price floor condition is satisfied. This is intuitive; larger w increases the profits that a firm can earn from the sale of add-ons, while smaller v reduces the monopolistic profits that a firm can enjoy by exploiting its captive consumers.

By ensuring that a deviation to monopolistic prices is unprofitable, the binding price floor condition describes the behaviour of firms at an add-on pricing equilibrium, if one exists: firms set $p_b = 0$, $p_a = w$ and share the contestable consumers while serving the captive consumers. Standard undercutting arguments do not permit equilibria at higher p_b. Each firm, thus, earns $\Pi^{add} = w\left(sq(1-q) + \frac{sq^2}{2} + \frac{(1-s)q}{2}\right)$.

However, while necessary for the existence of the add-on pricing equilibrium, the binding price floor condition cannot ensure that neither firm has an incentive to deviate to the all-inclusive strategy and undercut. Condition $\frac{1}{1+q} > s > \frac{1}{3(1-q)}$ guarantees that such deviations are also unprofitable.

To get an intuition behind why undercutting is unprofitable if $\frac{1}{1+q} > s > \frac{1}{3(1-q)}$, first notice that, by setting an add-on price w, an add-on firm charges an expected total price w to informed consumers (who always buy the add-on), and an expected total price wq to uninformed consumers (who buy the add-on with probability q). In other words, due to the voluntary nature of the add-on purchase, the add-on pricing strategy permits firms to price discriminate between informed and uninformed consumers. By following the all-inclusive strategy, the deviating firm cannot price discriminate; it offers the same deal to both consumer types. As such, it can undercut in one of two ways.

The first way is by setting some p_I slightly below w and stealing all the informed contestable consumers. At this p_I, the deviating firm cannot attract any uninformed consumers, as these consumers are better off purchasing from the add-on firm, at an expected total price wq. Hence, the deviating firm's profit is approximately $w(sq(1-q)+sq^2)$. Comparison with $w\left(sq(1-q)+\frac{sq^2}{2}+\frac{(1-s)q}{2}\right)$ yields that the deviation is not profitable if $\frac{1}{1+q} > s$. Intuitively, the fact that the deviating firm serves only informed consumers makes this deviation unprofitable if the proportion of uninformed consumers is sufficiently large.

The second way is by setting some p_I slightly below wq and stealing all contestable consumers, earning approximately $wq\big(sq(1-q)+sq^2+(1-s)\big)$. Comparison with $w\left(sq(1-q)+\frac{sq^2}{2}+\frac{(1-s)q}{2}\right)$ yields that the deviation is unprofitable if $s > \frac{1}{3(1-q)}$. Intuitively, this is because by setting a price that attracts uninformed consumers (wq), the deviating firm makes a discrete cut in the price paid by informed consumers, who otherwise paid w. This discrete price cut makes the deviation unprofitable if the proportion of informed consumers is large enough.

To see that a pure strategy equilibrium does not exist when either of these deviations is profitable, notice that a profitable deviation to an all-inclusive equilibrium implies that undercutting remains a best response to prices that yield positive mark-ups. Because of this, each firm's optimal behaviour is similar to that described in the non-binding price floor case: firms have an incentive to undercut when prices are relatively high, and to deviate to monopolistic prices when prices are sufficiently low.

2.4.2.1 Comparative statics

Given a binding price floor, higher matching quality q reduces the parameter space for which an add-on pricing equilibrium exists, as it increases the profits from both deviations.

The profits of a deviating firm that aims to steal all the informed contestable consumers increase because higher q increases the number of informed contestable consumers that the deviating firm can serve (sq^2) more than the number of contestable consumers that each firm serves at the add-on pricing equilibrium ($\frac{sq^2}{2} + \frac{(1-s)q}{2}$). The profits of a deviating firm that aims to steal all contestable consumers increase because higher q permits the deviating firm to undercut at a higher price (higher wq).

The binding price floor condition is more easily binding the larger the population of contestable consumers in the market ($sq^2 + (1-s)q$), and the smaller the population of captive consumers in the market ($sq(1-q)$). Whether a higher proportion of informed consumers (s) increases or decreases the parameter space for which the binding price floor is binding depends on how s affects contestable and captive consumer population. The binding price floor condition is more easily satisfied if the newly informed consumers are more likely to be contestable (probability q^2) and less easily satisfied if the newly informed consumers are more likely to be captives (probability $q(1-q)$). In a similar way, whether a higher match quality (q) is conducive to a binding price floor constraint depends largely on whether or not the higher q increases the population of captive consumers versus the population of contestable consumers.

Note that equilibrium condition $s > \frac{1}{3(1-q)}$ implies that $w > v - w$ is necessary for the price floor to be binding. Put differently, the value of the add-on must be greater than the value of the base good. To see this, rearrange (BF) to $w > (v-w)\frac{2sq(1-q)}{sq^2+(1-s)q}$ and notice that the fraction in the right-hand side is greater than 1 when $s > \frac{1}{3(1-q)}$. This suggests that an add-on pricing equilibrium with rational consumers and matching uncertainty requires that the product's functionality more than doubles after the purchase of the add-on. Dropbox is an example of a product that appears to have this feature; the basic (free) version offers 2 gigabytes of storage space, while paid version Dropbox Plus offers 1 terabyte of storage space and various additional services. To heavy users, 1 terabyte of storage space could be worth more than twice as much as the basic 2 gigabytes space.

If add-on valuation differs across consumers then $w > (v-w)\frac{2sq(1-q)}{sq^2+(1-s)q}$ only needs to hold on average. To see this, consider the following modification to the main model. Suppose that, among consumers that are correctly matched, the add-on has intrinsic value ω for a proportion α of consumers, and intrinsic value 0 for proportion $1-\alpha$. Suppose also that consumers and firms know the sizes of α and ω, but do not know in advance whether a consumer values the add-on or not. Then we can reinterpret w as the expected value of the add-on; that is, we can define w as $w=\alpha\omega$. Given monopoly power over the add-on, an add-on firm sets $p_a=\omega$ and earns, on average, $\alpha\omega$ from all its customers that are correctly matched. Applying $w=\alpha\omega$ to (BF) and rearranging gives

$$\alpha\omega > (v-\alpha\omega)\frac{2sq(1-q)}{sq^2+(1-s)q} \qquad (BF')$$

(BF') shows that a binding price floor requires that only the expected add-on value, $\alpha\omega$, is large enough. The price floor can be binding even if the add-on has no value for a proportion $1-\alpha$ of the consumer population, as long as ω is sufficiently high for the rest of the consumers. To the extent that ω may have extremely large values for certain consumers, (BF') makes our model consistent with evidence that the freemium business model relies on the presence of a few "whale" customers that spend large amounts of money on microtransactions.[25]

2.4.2.2 *Welfare*

Compared to a world in which firms cannot separate the base good from the add-on component (and, thus, can only follow the all-inclusive strategy), the add-on pricing equilibrium strictly improves both firm profits and social surplus.

The logic is as follows. If firms are constrained to playing the all-inclusive strategy (the *all-integrated game*), then the unique equilibrium is in mixed strategies.[26] Each firms chooses a price randomly from a support $[v, v(1-q)]$ if $q \leq \frac{2s-1}{2s}$, or

[25] https://www.theguardian.com/technology/appsblog/2011/jul/26/freemium-mobile-games-whales [accessed September 2018].

[26] For a detailed analysis of the mixed-strategy equilibrium when both firms follow the all-inclusive strategy see section 2.3 in Armstrong (2015). In his model, match quality is denoted by α and the population of informed consumers (which Armstrong refers to as "savvy") is denoted by σ.

$\left[v, v\frac{1-s+sq^2}{1-s+sq}\right] \cup [vq, v(1-q)]$ otherwise. Intuitively, firms have an incentive to compete on p_I, but each individual firm has an incentive to set $p_I = v$ and serve all its captive consumers if the rival's price is too low. Since every price in the equilibrium distribution must yield the same expected profits, equilibrium firms profits are given by $vsq(1-q)$; the profits that each firm earns by setting $p_I = v$ and serving only its captive consumers with probability 1. Since the price floor is binding at the add-on pricing equilibrium, it follows from condition (BF) that firm profits at the add-on pricing equilibrium are higher than $vsq(1-q)$.

To see why social surplus is higher at the add-on pricing equilibrium, notice that the equilibrium price support at the all-integrated game includes prices that are above the uninformed consumers' reservation value, vq. In particular, if $q \leq \frac{2s-1}{2s}$ the uninformed consumers are not served at any price in equilibrium, as $v(1-q) > vq$. Therefore, at a market with only integrated goods, uninformed consumers may be left unserved. In contrast, firms in the add-on pricing equilibrium always serve uninformed consumers.

2.5 Conclusion

This chapter aims to provide a rational explanation to why firms may adopt add-on pricing strategies in equilibrium when the adoption of add-on pricing constrains own firms' ability to offer competitive deals. We highlight two distinct mechanisms that can play a role in competitive firms' choice to unbundle base goods from add-ons. The first mechanism, relevant to a wide range of markets with physical add-on services (e.g. hotels, airlines) is that add-on pricing firms can save on costs if add-ons are not purchased. The second mechanism, particularly relevant to software markets, is that add-on pricing can permit firms to price discriminate between consumers that are informed about match quality and consumers that are uninformed. This type of price discrimination is not possible if firms bundle base goods and add-ons.

We show conditions under which either mechanism, on its own, is sufficient to establish an add-on pricing equilibrium even when consumers are fully rational and add-on pricing limits price-cutting. In the cost-savings model we find that an add-on pricing equilibrium emerges when the cost savings associated with the add-on pricing business strategy are sufficiently large. This holds even if the add-on pricing strategy is associated with a higher expected total price. In the matching uncertainty model, we find that an add-on pricing equilibrium emerges if the proportion of informed consumers is of moderate size and the add-on component is of sufficiently high value. For markets described by this model, add-on pricing is the unique equilibrium market outcome in pure strategies.

By unveiling non-behavioural mechanisms that may be responsible for the adoption of add-on pricing strategies in various markets in which price cutting is constrained, our work is informative to regulators in two ways. First, it shows that the adoption of add-on pricing may not be a supply-side response to consumers' behavioural biases. As such, demand-side market remedies that aim to improve consumers' decision making (e.g. by disclosing features or nudging consumers) may be ineffective. Second, our work shows that add-on pricing may in fact improve social surplus if it is an equilibrium strategy in markets with rational consumers. Thus, steering firms away from this business practise may be harmful.

In attempting to most clearly illustrate the role of cost savings and matching uncertainty, our models pay little attention to various other parameters that may be important in the firms' choice to unbundle. For instance, by assuming that products are perfect substitutes we overlook the role that competition intensity may play. The answer appears not to be straightforward, as firms offer add-ons in both highly competitive (e.g. apps) and very differentiated (e.g. airlines) markets. Another aspect that our models overlook is the importance of intermediaries, which are often used for the sale of products with add-ons (e.g. apps are sold in Play Store and Apple Store; hotel rooms are sold via platforms like Booking.com). A richer model may be able to study how the presence of a binding price floor interacts with firms' vertical agreements and provide conditions under which add-on pricing may arise.

2.6 Appendix A

Proof of Proposition 1.

Part (i): Non-binding price floor.

Stage 2

Suppose that $R_{i,j} = I$ in stage 1. Then in stage 2 firms compete on the all-inclusive price, p_I. Standard Bertrand arguments suggest that the unique equilibrium in stage 2 is characterised by $p_I^* = c$ and profits $\Pi^{inc} = 0$.

Now suppose that $R_{i,j} = A$ in stage 1. Then in stage 2 there exists a unique equilibrium in which both firms play $p_b^* = c - wq$ and each firm earns $\Pi^{add} = 0$. To see that p_b^* constitutes an equilibrium base price, notice that a firm that deviates to some $p_b' > p_b^*$ serves no consumers, while a firm that deviates to some $p_b' < p_b^*$ serves each consumer but makes losses. Thus p_b^* indeed characterises an equilibrium. Uniqueness follows from standard undercutting arguments.

Next, suppose that $R_i = I$ and $R_j = A$ in stage 1. Then in stage 2 there exists a unique equilibrium in which the add-on firm sets p_b^* such that $p_b^* + wq = c$, and the all-inclusive firm sets $p_I^* = c$. At this equilibrium, each firm earns $\Pi^{asym} = 0$. To see that p_b^* and p_I^* characterise an equilibrium, notice that any firm that deviates to a higher price serves no consumers, while any firm that deviates to a lower price serves each consumer but makes losses. Uniqueness follows from standard undercutting arguments.

Stage 1

Since $\Pi^{inc} = \Pi^{add} = \Pi^{asym}$, each firm is indifferent between the add-on pricing strategy and the all-inclusive strategy in stage 1. Therefore, both an add-on pricing equilibrium, and an all-inclusive equilibrium exist.

Part (ii): Binding price floor.

Stage 2

Suppose that $R_{i,j} = I$ in stage 1. Then in stage 2 firms compete on the all-inclusive price, p_I. Since the $f \leq c$, the price floor constraint is never binding. Therefore, standard Bertrand arguments suggest that a unique equilibrium in stage 2 is characterised by $p_I^* = c$ and profits $\Pi^{inc} = 0$.

Now suppose that $R_{i,j} = A$ in stage 1. Then in stage 2 there exists a unique equilibrium in which both firms play $p_b^* = f$ and each firm earns $\Pi^{add} = \frac{f+wq-c}{2}$. To see that p_b^* constitutes an equilibrium base price, notice that a firm that deviates to some $p_b' > p_b^*$ serves no consumers, while the price floor constraint does not allow firms to set any $p_b' < p_b^*$. Thus ${p_b^*}^b$ indeed characterises an equilibrium. Uniqueness follows from standard undercutting arguments.

Next, suppose that $R_i = I$ and $R_j = A$ in stage 1. Then in stage 2 there exists a unique equilibrium in which the add-on firms sets $p_b^* = f$, which corresponds to an expected total price $f + wq$, and the all-inclusive firm sets some $p_I^* = f + wq - \epsilon$, where ϵ is infinitesimal. Since $p_I^* < f + wq$, the all-inclusive firm serves every consumer, earning $\Pi_I^{asym} = f + wq - c$, while the add-on pricing firm serves no consumers and earns $\Pi_A^{asym} = 0$. Uniqueness follows from standard undercutting arguments.

Stage 1

Since $\Pi_I^{asym} > \Pi^{add}$, each firm is better off choosing the strategy I when its rival chooses strategy A. Therefore, an add-on pricing equilibrium does not exist. Since $\Pi_A^{asym} = \Pi^{add}$, each firm is indifferent between strategies A and I, when its rival chooses strategy I. Therefore, an all-inclusive equilibrium exists.■

Proof of Proposition 2.

Part (i): Non-binding price floor.

Stage 2

Suppose that $R_{i,j} = I$ in stage 1. Then in stage 2 each firm sets $p_I^* = c$ and earns profits $\Pi^{inc} = 0$. For proof see part *(i)*, in the proof of Proposition 1 when $R_{i,j} = I$.

Now suppose that $R_{i,j} = A$ in stage 1. Then in stage 2 each firm sets p_b^* such that $p_b^* + wq = c_b + c_a q$ and earns profits $\Pi^{add} = 0$. For proof replace $c_b + c_a q$ with c and follow part *(i)* in the proof of Proposition 1 when $R_{i,j} = A$.

Next, suppose that $R_i = I$ and $R_j = A$ in stage 1. Then in stage 2, if $c > c_b + c_a q$ there exists a unique equilibrium in which the all-inclusive firm sets $p_I^* = c$, and the add-on firm sets p_b^* such that $p_b^* + wq = c - \epsilon$. At this equilibrium, the add-on pricing firm serves every consumer, earning profits $\Pi_A^{asym} = c - c_b - c_a q - \epsilon$. The all-inclusive firm does not serve any consumers, and, thus, earns $\Pi_I^{asym} = 0$. To see why the pair (p_I^*, p_b^*) constitutes an equilibrium, notice that the all-inclusive firm has no incentive to deviate to any $p_I^d \neq p_I^*$, as at $p_I^d > p_I^*$ the firm does not attract any customers, and at $p_I^d < p_I^*$ the firms serves all consumers at a loss. The add-on pricing firm has no incentive to deviate to $p_b^d = p_I^*$, as that results in profits $\frac{c - c_b - c_a q}{2} < \Pi_A^{asym}$. It also has no incentive to deviate to $p_b^d > p_I^*$, as that results in zero demand. Finally, it has no incentive to deviate to $p_b^d < p_b^*$, as that result only in a lower mark-up, with no effect on demand. Uniqueness follows from standard undercutting arguments.

If $c < c_b + c_a q$ there exists a unique equilibrium in which the add-on firm sets p_b^* such that $p_b^* + wq = c_b + c_a q$, and the all-inclusive firm sets $p_I^* = c_b + c_a q - \epsilon$. At this equilibrium, the all-inclusive firm serves every consumer, earning profits $\Pi_I^{asym} = c_b + c_a q - c - \epsilon$. The add-on firm does not serve any consumers, and, thus, earns $\Pi_A^{asym} = 0$. The arguments as to why the pair (p_I^*, p_b^*) constitutes an equilibrium are analogous to those for the case in which $c > c_b + c_a q$. Uniqueness follows from standard undercutting arguments.

If $c = c_b + c_a q$ there exists a unique equilibrium in which the add-on firm sets p_b^* such that $p_b^* + wq = c$, and the all-inclusive firm sets $p_I^* = c$. At this equilibrium, each firm earns $\Pi^{asym} = 0$. To see that p_b^* and p_I^* characterise an equilibrium, notice that any firm that deviates to a higher price serves no consumers, while any firm that deviates to a lower price serves each consumer but makes losses. Uniqueness follows from standard undercutting arguments.

Stage 1

Suppose $c > c_b + c_a q$. Then, comparing Π_I^{asym} with Π^{add} yields $\Pi_I^{asym} = \Pi^{add}$. Therefore, an add-on pricing equilibrium exists. Comparing Π_A^{asym} with Π^{inc} yields $\Pi_A^{asym} > \Pi^{inc}$. Therefore, an all-inclusive equilibrium does not exist. Thus if $c > c_b + c_a q$ an add-on pricing equilibrium exists and is the unique symmetric subgame perfect equilibrium in pure strategies.

Now suppose $c < c_b + c_a q$. Comparing Π_I^{asym} with Π^{add} yields $\Pi_I^{asym} > \Pi^{add}$. Therefore, an add-on pricing equilibrium does not exist. Comparing Π_A^{asym} with Π^{inc} yields $\Pi_A^{asym} = \Pi^{inc}$. Therefore, an all-inclusive equilibrium exists. Thus if $c > c_b + c_a q$ an all-inclusive equilibrium exists and is the unique symmetric subgame perfect equilibrium in pure strategies.

Finally, suppose $c = c_b + c_a q$. In this case $\Pi_I^{asym} = \Pi^{add}$ and $\Pi_A^{asym} = \Pi^{inc}$. Therefore, both an all-inclusive equilibrium and an add-on pricing equilibrium exist.

Part (ii): Binding price floor.

Stage 2

Suppose that $R_{i,j} = I$ in stage 1. Then in stage 2 there exists a unique equilibrium in which each firm sets $p_I^* = c$ and earns profits $\Pi^{inc} = 0$. For proof see part *(i)*, in the proof of Proposition 1 when $R_{i,j} = I$.

Now suppose that $R_{i,j} = A$ in stage 1. Then in stage 2 there exists a unique equilibrium in which both firms play $p_b^* = f$ and each firm earns $\Pi^{add} = \frac{f + wq - c_b - c_a q}{2}$. For proof replace $c_b + c_a q$ with c and follow part *(i)* in the proof of Proposition 1 when $R_{i,j} = A$.

Next, suppose that $R_i = I$ and $R_j = A$ in stage 1. Then in stage 2, if $c > f + wq$ there exists a unique equilibrium in which the all-inclusive firm sets $p_I^* = c$, and the add-on firm sets p_b^* such that $p_b^* + wq = c - \epsilon$. At this equilibrium, the add-on pricing firm serves every consumer, earning profits $\Pi_A^{asym} = c - c_b - c_a q - \epsilon$. The all-inclusive firm does not serve any consumers, and, thus, earns $\Pi_I^{asym} = 0$. To see why the pair (p_I^*, p_b^*) constitutes an equilibrium, notice that the all-inclusive firm has no incentive to deviate to any $p_I^d \neq p_I^*$, as at $p_I^d > p_I^*$ the firm does not attract any customers, and at $p_I^d < p_I^*$ the firms serves all consumers at a loss. The add-on pricing firm has no incentive to deviate to $p_b^d = p_I^*$, as that results in profits $\frac{c - c_b - c_a q}{2} < \Pi_A^{asym}$. It also has no incentive to deviate to $p_b^d > p_I^*$, as that results in zero demand. Finally, it has no incentive to deviate to $p_b^d < p_b^*$, as that result only in a lower mark-up, with no effect on demand. Uniqueness follows from standard undercutting arguments.

If $c < f + wq$ there exists a unique equilibrium in which the add-on firm sets $p_b^* = f$, which corresponds to an expected total price $f + wq$, and the all-inclusive firm sets $p_I^* = f + wq - \epsilon$. At this equilibrium, the all-inclusive firm serves every consumer, earning profits $\Pi_I^{asym} = f + wq - c - \epsilon$. The add-on firm does not serve any consumers, and, thus, earns $\Pi_A^{asym} = 0$. The arguments as to why the pair (p_I^*, p_b^*) constitutes an equilibrium are analogous to those for the case in which $c > f + wq$. Uniqueness follows from standard undercutting arguments.

If $c = f + wq$ then in the unique equilibrium the add-on pricing firm sets $p_b^* = f$ and the all-inclusive firm sets $p_I^* = c$. Consumers are shared evenly. The add-on pricing firm earns $\Pi_A^{asym} = \frac{f + wq - c_b - c_a q}{2}$, while the all-inclusive firm earns $\Pi_I^{asym} = 0$. To see why the pair (p_I^*, p_b^*) constitutes an equilibrium, notice that the all-inclusive firm has no incentive to deviate to any $p_I^d \neq p_I^*$, as at $p_I^d > p_I^*$ the firm does not attract any customers, and at $p_I^d < p_I^*$ the firms serves all consumers at a loss. The add-on pricing firm has no incentive to deviate to $p_b^d > p_b^*$, as that results in zero demand, and the price floor does not permit a deviation at $p_b^d < p_b^*$. Uniqueness follows from standard undercutting arguments.

Stage 1

Suppose $c > f + wq$. Then, comparing Π_I^{asym} with Π^{add} yields $\Pi^{add} > \Pi_I^{asym}$. Therefore, an add-on pricing equilibrium exists. Comparing Π_A^{asym} with Π^{inc} yields $\Pi_A^{asym} > \Pi^{inc}$. Therefore, an all-inclusive equilibrium does not exist. Thus if $c > f + wq$ an add-on pricing equilibrium exists and is the unique symmetric subgame perfect equilibrium in pure strategies.

Now suppose $c < f + wq$. Comparing Π_I^{asym} with Π^{add} yields $\Pi^{add} \geq \Pi_I^{asym}$ if $c - c_b - c_a q \geq f + wq - c$. Therefore, if $c - c_b - c_a q \geq f + wq - c$ an add-on pricing equilibrium exists. If $c - c_b - c_a q < f + wq - c$ each firm prefer to choose strategy I when its rival chooses strategy A; that is, an add-on pricing equilibrium does not exist. Comparing Π_A^{asym} with Π^{inc} yields $\Pi_A^{asym} = \Pi^{inc}$. Therefore, if $c > f + wq$ an all-inclusive equilibrium exists.

Finally, suppose $c = f + wq$. Comparing Π_I^{asym} with Π^{add} yields $\Pi^{add} > \Pi_I^{asym}$. Therefore, an add-on pricing equilibrium exists. Comparing Π_A^{asym} with Π^{inc} yields $\Pi_A^{asym} > \Pi^{inc}$. Therefore, an all-inclusive equilibrium does not exist. Thus if $c = f + wq$ an add-on pricing equilibrium exists and is the unique symmetric subgame perfect equilibrium in pure strategies. ■

Proof of Proposition 3.

Part (i): Non-binding price floor.

Suppose that both firms choose strategy I and set $p_I = 0$, each earning zero profits. Then firm j can increases its profits by choosing strategy I and setting $p_I = v$. At this price it sells to all its captive consumers, earning $vsq(1 - q)$.

Now suppose that both firms choose strategy A and set $p_b = 0$. At this price each firm earns $\Pi^{floor} = w\left(sq(1 - q) + \frac{sq^2}{2} + \frac{(1-s)q}{2}\right)$. Then, again, firm j can increases its profits by choosing strategy I and setting $p_I = v$. At this price it sells to all its captive consumers, earning $vsq(1 - q)$. By the condition in part *(i)*, this deviation is always profitable.

Standard undercutting arguments suggest that any $p_I > 0$ and any $p_b > 0$ is not an equilibrium, as a rival can undercut and steal all contestable consumers. Therefore, a symmetric pure strategy price equilibrium does not exist.

Part (ii): Binding price floor.

Suppose that both firms choose strategy A and set $p_b = 0$. At this price each firm earns profits $\Pi^{floor} = w\left(sq(1-q) + \frac{sq^2}{2} + \frac{(1-s)q}{2}\right)$. The price floor constraint does not allow any firm to set $p_b < 0$. We, therefore, are left with two possible deviations:

(a) A deviation in which firm j chooses strategy I and undercuts at a price that steals all informed consumers. The highest price that achieves this is some $p_I = w - \epsilon$. At this price, the deviating firm earns, approximately, $w(sq(1-q) + sq^2)$. Comparison with Π^{floor} yields that the deviation is not profitable if $\frac{1}{1+q} > s$.

(b) A deviation in which firm j chooses strategy I and undercuts at a price that steal all consumers. The highest price that achieves this is some $p_I = wq - \epsilon$. At this price, the deviating firm earns, approximately, $wq(sq(1-q) + sq^2 + 1 - s)$. Comparison with Π^{floor} yields that the deviation is not profitable if $s > \frac{1}{3(1-q)}$.

Therefore, if $\frac{1}{1+q} > s > \frac{1}{3(1-q)}$ there exists an add-on pricing equilibrium. Uniqueness follows from part *(i)* of this proof.■

Chapter 3

Paying for Unwanted Add-ons in a Market with Endogenous Product Differentiation

3.1 Introduction

Credit card users forget payments and incur late payment charges. Mobile phone users accidentally incur roaming charges and experience bill-shock. Insurance buyers are surprised by unexpected excesses. Receivers of free trials pay unanticipated monthly charges for unwanted membership plans. What these examples show is that consumers often make mistakes and face unexpected costs. Some of these costs are the result of unforeseen contingencies. Others arise due to consumers' limited memory, their insufficient attention to product details, or other behavioural biases (Grubb, 2015a).

While some consumers are careful, others fail to avoid mistakes. Firms adjust their business strategies to take advantage of this. For example, firms may offer back-loaded contracts (e.g. DellaVigna and Malmendier, 2004; Grubb, 2015b), draw consumers' attention to the product's best attributes (Bordalo *et al.*, 2016), shroud certain price components (Gabaix and Laibson, 2006) and invest in the development of more innovative traps (Heidhues *et al.*, 2016).

In this chapter we show that the firms' incentives to profitably exploit consumer mistakes may also affect their choice of horizontal product differentiation.[27] Commonly, a firm's choice of product differentiation is determined by two forces: (i) the incentive to soften competition by offering a more differentiated product (the price effect), and (ii) the incentive to capture more consumers from rivals by offering a less differentiated product (the demand effect). We explore how the presence of naïve consumers interacts with these forces.

In order to more simply capture the problem of firms exploiting consumer mistakes, we model a market in which products feature a salient upfront price and a fee that is associated with the purchase of an "unwanted" add-on service. Such an unwanted add-on may be a roaming service used by accident, or an unconscious enrolment to a costly membership plan. Viewed more broadly, the cost of an unwanted add-on may represent the late payment cost incurred by a forgetful credit card user, or the unexpectedly limited coverage of an insurance plan. "Sophisticated" consumers avoid making mistakes and only pay the upfront price, while "naïve" consumers accidentally purchase the unwanted add-ons.

Standard competition arguments suggest that firms cannot profit from selling add-ons to naïve consumers. The prospect of high future revenues incentivises firms to compete more fiercely, lowering upfront prices possibly below marginal cost and dissipating profits from the add-on fees.[28] However, this stronger competitive pressure may not always succeed at competing away all revenues from add-on fees. Constraints may be limiting the extent to which firms can compete on upfront prices, allowing firms to retain some excess profits. To capture the idea that competition does not

---

[27] Horizontal differentiation may be real or spurious (Spiegler, 2006; Tremblay and Polasky, 2002). We model it in a traditional way, as a choice of product characteristics over which consumers have heterogeneous preferences.

[28] Gabaix and Laibson (2006) contains such a profit irrelevance result regarding unanticipated valuable add-ons. A similar dissipation argument appears in models where firms have monopoly power over an aftermarket due to switching costs (see, for example, Lal and Matutes, 1994; Shapiro, 1995).

always erode exploitation profits completely, we assume that upfront prices are bound below by a price floor.[29]

In real markets, a price floor may arise if price cutting attracts a disproportionate amount of unprofitable consumers (Ellison, 2005). For example, setting upfront prices below zero may attract a large proportion of sophisticated types that avoid any additional prices (Heidhues *et al.*, 2017). The presence of a zero price floor is hinted by the zero annual fees in numerous consumer credit card products (DellaVigna and Malmendier, 2004) and the free-if-in-credit business model for personal current accounts (Armstrong and Vickers, 2012). For products with aftermarket services, a price floor may be established by the consumers' ability to buy multiple units of the base good (Miao, 2010; Michel, 2017). Price floors may also arise due to market regulations (below-cost pricing may be viewed as prohibited predatory behaviour; regulation in the mutual funds industry may rule out "signing bonuses", as modelled in Heidhues *et al.* (2017)), due to consumers' perception of high upfront prices as a signal for quality (Wolinsky, 1983), or if consumers are suspicious of extremely low prices, expecting a trap.

We model firms' endogenous choice of horizontal differentiation as a choice of locations on a Hotelling line. The degree of horizontal differentiation is crucial to whether firms can enjoy any excess profits. Firms face two conflicting incentives regarding their location choice: (i) locate far apart to soften competition and enjoy higher prices, or (ii) locate close together to capture more consumers from rivals, while retaining some profits from the add-on fees through the price floor. Our model shows the conditions that influence the balance of these forces.

In contrast with traditional Hotelling models (Hotelling, 1929; d'Aspremont *et al.*, 1979; Neven, 1985) in which fixed prices lead to minimum differentiation and variable prices lead to maximum differentiation, we show that the firms' ability to exploit consumer mistakes may establish a maximum or a minimum-differentiation equilibrium, depending on market parameters. Firms may be better off at the

[29] See sections 2.2 and 2.3 in Heidhues and Koszegi (2018) for a review of articles in which price floors matter.

maximum-differentiation equilibrium despite dissipating all revenues from add-on fees.

The impact of consumer exploitation on product differentiation raises particular issues for consumer protection policy. The ex-ante degree of product differentiation, as well as the possibility that firms respond to interventions by choosing different product locations influence the effectiveness of conventional regulatory policies (e.g. informing consumers about the existence of costly add-ons, putting a cap on add-on fees).

Our model contributes to the growing literature in Behavioural Industrial Organisation (Heidhues and Koszegi, 2018; Grubb, 2015a) by investigating how the presence of behavioural consumers may affect firms' choice of product differentiation. In that our work is most closely related to Schultz (2004), who reveals that a more opaque price strengthens the competition softening effect of product differentiation, incentivising firms to offer more differentiated products. We, instead, show that a combination of a transparent and an opaque price component can induce firms to compete over product locations, resulting in less differentiated products.

To better distinguish the effect of consumer exploitation on equilibrium product differentiation, we abstract away from other issues that may arise when consumers incur unexpected fees. In assuming unit demands we exclude the possibility that the subsidised upfront prices may result in excessive consumption of the primary service (Heidhues and Koszegi, 2010; 2015). We develop a model in which sophisticated consumers can costlessly avoid any unexpected fees. However, avoidance may require pecuniary or effort costs, or an inefficient altering of consumption decisions (Gabaix and Laibson, 2006; Grubb, 2015b). In modelling an exogenous proportion of naïve consumers, we abstract away from the question whether firms (Gabaix and Laibson, 2006; Heidhues *et al*. 2017; Wenzel, 2014) or intermediaries (Murooka, 2015) have an incentive to educate such consumers.

Our model shows that consumer exploitation can have a knock-on effect on product differentiation, which may be beneficial to welfare. Other knock-on effects may not be as benign. Michel (2018), Gamp and Krahmer (2018) and Heidhues *et al.* (2017) reveal that the possibility to exploit naïve consumers can establish markets with products of inefficiently low quality. Heidhues *et al.* (2016) show that profitable exploitation can incentivise firms to invest towards more exploitative contracts, taking resources away from value-enhancing R&D.

We contribute to the academic discussion regarding whether the presence of sophisticated consumers benefits (e.g. Varian, 1980; Armstrong and Chen, 2009; Heidhues and Koszegi, 2017) or harms (e.g. Gabaix and Laibson, 2006; Armstrong and Vickers, 2012) naïve consumers. In particular, our model exhibits what Armstrong (2015) dubs a "rip-off externality": the add-on expenses of naïve consumers cross-subsidise the spending of sophisticated types. Going further, we show that, given endogenous product differentiation, the harm to naïve consumers may go beyond the first-order effect of the cross-subsidisation. The presence of a sufficiently large proportion of sophisticated consumers may establish a maximum-differentiation equilibrium, resulting in disproportionately higher upfront prices. Viewed the other way around, naïve types may benefit sophisticates not only directly, through funding lower upfront prices, but also indirectly, by facilitating a minimum-differentiation equilibrium.

Finally, our model warns of previously understudied supply-side responses that may undermine the effectiveness of conventional regulatory policies, or even lead to more adverse market outcomes. In studying unintended effects of regulatory interventions, Grubb (2015b) warns that bill-shock regulation can harm social welfare when firms offer multiple contracts to discriminate between low- and high-demand consumers. This is because bill-shock regulation forces firms to introduce quantity distortions in order to discriminate. Armstrong *et al.* (2009) show that price controls dampen consumers' incentives to become informed, which may lead to higher prices. Ellison and Wolitzky (2012) demonstrate that firms may cancel out the beneficial effects of disclosure policies by increasing their obfuscation efforts.

The rest of the paper is organised as follows. Section 3.2 sets up our model. Section 3.3 presents our equilibrium analysis. In section 3.3.1 we characterise the price equilibrium for fixed symmetric locations. In section 3.3.2 we characterise the location stage given equilibrium prices. In Section 3.4 we study welfare and policy implications. Section 3.5 concludes.

3.2 Model

Consider two firms located at l_1, l_2 on a Hotelling line between 0 and 1, and, without loss of generality, let $l_1 \leq l_2$. Each firm produces a homogenous good of utility v at cost c, $c < v$. Assume that v is large enough that the market is always fully covered.

Consumers are uniformly distributed on the interval $[0,1]$ and have mass 1. Each consumer has unit demand, and may visit and purchase from at most one firm. For a consumer located at $\chi \in [0,1]$, visiting and purchasing from firm i entails a distance cost $t(\chi - l_i)^2$.

Purchasing product i entails two price components: an upfront price $f_i \in \left[\underline{f}, +\infty\right)$, where $\underline{f} < c$ represents an exogenous price floor, and an exogenous add-on fee $\bar{\bar{a}}$. Purchase of product i requires payment of both prices, unless the add-on fee is avoided. Avoidance is costless.[30] Consumers have different types depending on whether they are aware of the add-on fee. Sophisticated consumers (proportion $q \in [0,1]$) are aware the add-on fee and costlessly avoid it. Naïve consumers (proportion $1 - q$) are unaware before making a purchase choice and pay $\bar{\bar{a}}$ upon purchase.[31] Both types are uniformly distributed on the line.

[30] Costly avoidance is associated with a welfare loss (e.g. Gabaix and Laibson, 2006) whose study is not central to our model. A positive but relatively small avoidance cost $\bar{a}$ does not alter our results qualitatively, as long as it satisfies $\bar{a} \leq \bar{\bar{a}}(1 - q)$. Chapter 4 studies a model in which an avoidance cost exists and is high enough to violate this inequality.

[31] The add-on fee represents a contingent charge that consumers may incur by mistake (e.g. by entering an unauthorised overdraft). In this case $\bar{\bar{a}}$ represents the maximum charge permitted by the regulator. Our model assumes that such charges are only incurred accidentally, so firms have no incentive to set a charge below $\bar{\bar{a}}$. In a more dynamic setting, the add-on fee could represent a periodic charge for having access to a service after the trial period. In a world in which consumers are paying for such service primarily by accident, $\bar{\bar{a}}$ represents the price above which naïve consumers become aware of the recurring expenses and avoid any further payments. Finally, the add-on fee may represent the money that disengaged back-book customers leave on the table by not switching to a different product. In that case, $\bar{\bar{a}}$ represents the price below which back-book customers remain disengaged.

The timing is as follows. In stage 1 firms simultaneously choose locations on the line. In stage 2 firm simultaneously set upfront prices. Consumers observe firm locations and upfront prices, form expectations about whether add-on fees apply, and choose from which firm to purchase. We solve the game using backward induction looking for symmetric equilibria.

3.3 Equilibrium analysis

3.3.1 Equilibrium price structure for fixed symmetric locations

We first analyse the Nash equilibrium in prices for any given symmetric location pair (l_1, l_2). With full coverage, each firms' demand is determined by the position of the consumer who is indifferent between buying from firm i versus buying from firm j. The add-on fees do not affect consumers' purchase decision, as sophisticated types avoid them while naïve types ignore them. Hence, the location of the indifferent consumer is given by x, which satisfies $f_1 + t(x - l_1)^2 = f_2 + t(x - l_2)^2$. Solving for x gives that $x = \frac{l_1 + l_2}{2} - \frac{f_1 - f_2}{2t(l_2 - l_1)}$. At any symmetric pure strategy price equilibrium each firm serves at least some consumers, so the indifferent consumer in located between 0 and 1. Hence, for prices around an equilibrium upfront price f^*, x represents the demand of firm 1 and $1 - x$ represents the demand of firm 2.

The add-on fees are only paid by the naïve types. Thus, firm 1's maximisation problem is $\max_{f_1} \Pi_1 \; s.t. f_1 \geq \underline{f}$, where $\Pi_1 = (f_1 + \bar{\bar{a}}(1 - q) - c)x$. Similarly, firm 2's maximisation problem is $\max_{f_2} \Pi_2 \; s.t. \; f_2 \geq \underline{f}$, where $\Pi_2 = (f_2 + \bar{\bar{a}}(1 - q) - c)(1 - x)$. Profit maximisation yields the following best response functions:

$$f_1^{BR} = \max\left(\frac{t(l_1 + l_2)(l_2 - l_1) + c - \bar{\bar{a}}(1 - q) + f_2}{2}, \underline{f}\right)$$

$$f_2^{BR} = \max\left(\frac{t(2 - l_1 - l_2)(l_2 - l_1) + c - \bar{\bar{a}}(1 - q) + f_1}{2}, \underline{f}\right)$$

The slope of the best response functions guarantees that there exists a unique and stable[32] price equilibrium for each location pair. Depending on firm locations, this equilibrium may be characterised by a binding price floor constraint for none, one, or both firms, as best response functions are kinked at $\underline{f}$. To gain better intuition we focus on the equilibrium outcome when firm locations are symmetric.[33]

3.3.1.1 Prices

Proposition 1 presents the price equilibrium for symmetric locations ($l_1 = 1 - l_2$). Subscripts in f^*, a^*denote the range of l_2 for which these variables characterise a pure strategy price equilibrium. For example, if $l_2 \in l_2^H$ the pure strategy price equilibrium is characterised by f_H^*, a_H^*, and Π_H^*.

Proposition 1. (Price equilibria for symmetric fixed locations). *Let* $l_1 = 1 - l_2$.

(a) *For* $l_2 \in l_2^H = \left[\frac{1}{2} + \frac{\underline{f} + \bar{\bar{a}}(1-q) - c}{2t}, 1\right]$ *there exists a pure strategy price equilibrium with* $f_H^* = c - \bar{\bar{a}}(1-q) + t(2l_2 - 1)$ *and* $\Pi_H^* = \frac{t(2l_2-1)}{2}$.

(b) *For* $l_2 \in l_2^L = \left[\frac{1}{2}, \frac{1}{2} + \frac{\underline{f} + \bar{\bar{a}}(1-q) - c}{2t}\right)$ *there exists a pure strategy price equilibrium with* $f_L^* = \underline{f}$ *and* $\Pi_L^* = \frac{\underline{f} + \bar{\bar{a}}(1-q) - c}{2}$.

Proof: *In appendix B.*

Proposition 1 reveals how product differentiation, traditionally represented by $t(l_2 - l_1)$ in Hotelling models, affects price structure when a proportion of naïve

[32] Each best response function has slope $r_i = \frac{1}{2}$ at the unconstrained part and $r_i = 0$ at the constrained part, satisfying the stability condition $|r_i| < 1$ (Dixit, 1986).

[33] An analysis of the pricing stage for general locations is required for the derivation of Proposition 2 in section 3.3.2. For details see Appendix B.

consumers does not anticipate paying add-on fees. To better comprehend the effect of product differentiation, it is useful to decompose it into two distinct elements: inherent differentiability, denoted by t, and the degree to which firms exploit the inherent differentiability, denoted by interfirm distance $l_2 - l_1$. For ease in exposition we make use of the symmetry in firm locations and indicate interfirm distance by l_2, with high (low) l_2 representing large (small) interfirm distance.

With no incentive to lower add-on fees, firms can only compete on upfront prices. Proposition 1 reveals that the relationship between equilibrium upfront prices and interfirm distance is not smooth. For $l_2 \in l_2^H$ equilibrium upfront prices become smaller as firms locate closer together, since competition is more intense with lower interfirm distance. This is in accordance with the predictions of the standard endogenous prices Hotelling model (d'Aspremont *et al.*, 1979; Neven, 1985). For $l_2 \in l_2^L$ the relationship between equilibrium upfront prices and interfirm distance breaks down, as the price floor constraint does not permit upfront prices to be competed below $\underline{f}$.

The extent to which changes in interfirm distance correspond to different price structures depends crucially on the relative size of $\underline{f}$, c, $\bar{\bar{a}}$, q and t.

Representing the price floor, higher $\underline{f}$ increases the parameter space for which the price floor constraint ($f_i \geq \underline{f}$) is binding. As such, it increases the parameter space for which $l_2 \in l_2^L$ and reduces the parameter space for which $l_2 \in l_2^H$. A higher maximum add-on fee ($\bar{\bar{a}}$) and a lower proportion of sophisticated types (q) also increase the parameter space for which the price floor constraint is binding, leading to a larger l_2^L and a smaller l_2^H. To get an intuition recall that firms use revenues from add-on fees to subsidise upfront prices.[34] This implies the price floor constraint is more likely to be binding with higher $\bar{\bar{a}}$ and lower q, as the former increases the add-on fees that they earn from every naïve customer, and the latter increases the proportion of naïve types in the market.

[34] Subsidisation of upfront prices is common in problems with second-period prices (Shapiro, 1995; Gabaix and Laibson, 2006). The extent of the subsidisation in the context of our model is discussed further below.

A higher marginal cost (c) reduces the parameter space for which the price floor is binding, as it increases equilibrium upfront prices for any given location pair. This results in a smaller l_2^L and a larger l_2^H. The effect of inherent differentiability (t) on price structure moves in the same direction as the effect of c. Intuitively, higher t increases the extent to which higher interfirm distance affects upfront prices. As a result, it increases the parameter space for which the price floor is slack.

3.3.1.2 Profits

It is an established result in competitive markets that firms enjoy higher profits the more differentiated the products are. This does not necessarily apply in markets with naïve consumer types and unexpected add-on fees. While large increases in product differentiation increase industry profits, small increases may not. This result follows directly from Proposition 1, and is described in Corollary 1.

Corollary 1. (Equilibrium profits and interfirm distance)

(a) Π_H^* *is increasing in* l_2. Π_L^* *is independent of* l_2.

(b) *Firms prefer symmetric locations with* $l_2 \in l_2^H$ *over symmetric locations with* $l_2 \in l_2^L$.

Proof: *In appendix B.*

Markets shares are not affected by symmetric increases in interfirm distance. As such, any change in profits arises from changes in equilibrium upfront prices. Part *(a)* states that equilibrium profits increase with small increases in interfirm distance only if interfirm distance is originally high enough that the price floor constraint is slack. When the price floor constraint is binding ($l_2 \in l_2^L$) firms cannot translate higher interfirm distance into higher equilibrium profits, as upfront prices are fixed.

Part *(b)* of Corollary 1 states that, given changes in product differentiation large enough to change the equilibrium market structure, firms prefer high over low product differentiation. This follows from part *(a)*. To see why, first compare Π_H^* and Π_L^*, and see that $\Pi_H^* > \Pi_L^*$ for any $l_2 \in l_2^H$. Given that Π_H^* is increasing in l_2, while Π_L^* is independent of l_2, we can conclude that equilibrium profits are higher when $l_2 \in l_2^H$.

A crucial problem in markets with naïve consumers and add-on fees is understanding the extent to which firms enjoy extra profits from such fees; in other words, the extent to which revenues from such fees subsidise lower upfront prices. To better approach this question it is useful to derive the profits that firms would enjoy when naïve consumers are absent – the sophisticated-consumers benchmark. We define the benchmark as the result of each firm solving the unconstrained maximisation problem $\max_{f_i}(f_i - c)\left(\frac{1}{2} - \frac{f_i - f^*}{2t(2l_2-1)}\right)$.[35] Maximisation yields equilibrium profits at the sophisticated-consumers benchmark $\Pi_{bench}^* = \frac{t(2l_2-1)}{2}$.

The extent to which firms enjoy profits beyond the sophisticated-consumers benchmark depends largely on how differentiated the products are. Firms fully erode fee revenues when product differentiation is relatively high ($l_2 \in l_2^H$), as illustrated by the fact that f_H^* is reduced by $\bar{\bar{a}}(1-q)$.[36] Full subsidisation of upfront prices is possible due to the price floor constraint being slack even when subsidy $\bar{\bar{a}}(1-q)$ is taken into account. The complete erosion of fee revenues results in equilibrium profits that are identical to those in the sophisticated-consumers benchmark.[37]

In contrast, subsidisation is only partial for relatively low product differentiation ($l_2 \in l_2^L$), as the price floor constraint does not permit upfront prices to be subsidised

---

[35] This is the profit maximisation problem in the standard Hotelling model with variable prices and quadratic distance costs.

[36] In this case fee revenues are conceptually equivalent to a reduction in marginal costs by $\bar{\bar{a}}(1-q)$. In our model the pass-through rate to such implicit cost reduction is 1. For more general settings that host downward-sloping demands and variable marginal costs, pass-through rate depends on demand and supply elasticities. For a general treatment of cost pass-through, see Weyl and Fabinger (2013). For an application of cost pass-through analysis to markets with hidden fees, see Agarwal *et al*. (2014).

[37] Note that, due to the subsidisation, the equilibrium upfront price may even be below marginal cost. This is consistent with the loss leader business model that characterises various markets with add-on services (e.g. retail financial products). For an early model of loss-leaders see Lal and Matutes (1994).

below c. As a result, for given locations, firm profits are higher than those in the sophisticated-consumers benchmark. To see this compare Π_L^* with Π_{bench}^* and see that $\Pi_L^* > \Pi_{bench}^*$ for $l_2 < \frac{1}{2} + \frac{\underline{f} + \bar{\bar{a}}(1-q) - c}{2t}$; that is, for $l_2 \in l_2^L$.

Figure 3.1 illustrates the result of Corollary 1, depicting how profits of firm 2 change with the degree of product differentiation.[38] Profits of firm 1 follow by symmetry. To ensure existence of location pairs that satisfy $l_2 \in l_2^L$ and existence of location pairs that satisfy $l_2 \in l_2^H$, figure 3.1 assumes $t > \underline{f} + \bar{\bar{a}}(1-q) - c$. The solid line represents the profit functions described in Proposition 1. Profits are strictly increasing in l_2 for $l_2 \in l_2^H$ and constant for $l_2 \in l_2^L$. Furthermore, profits for $l_2 \in l_2^H$ are higher than profits for $l_2 \in l_2^L$. The dotted line represents profits at the sophisticated-consumers benchmark, strictly increasing in l_2 regardless of interfirm distance. The vertical distance between the solid and the dotted line represents the additional surplus that firm 2 enjoys due to retained add-on revenues – the add-on revenues that remain after subsidisation of upfront prices has been exhausted by the binding price floor constraint.

3.3.1.3 Consumer effects

Being the only type that pays the add-on fee, naïve consumers do not enjoy the full extent of the subsidy on upfront prices. Sophisticated consumers also benefit from the lower upfront prices, as firms cannot offer a discount solely to naïve types. As a result, naïve types pay more for the same product in equilibrium. In particular, for $l_2 \in l_2^H$ naïve types pay total price $p_n^H = c + \bar{\bar{a}}q + t(2l_2 - 1)$, while sophisticated types pay only $p_s^H = c - \bar{\bar{a}}(1-q) + t(2l_2 - 1)$. For $l_2 \in l_2^L$ naïve types pay total price $p_n^L = \underline{f} + \bar{\bar{a}}$, while sophisticated types pay only $p_s^L = \underline{f}$. This creates a distribution issue, which can be particularly worrying if naïve consumer behaviour is associated with vulnerable consumers (Financial Conduct Authority, 2015).

[38] Note the similarities between this figure and figure 4.1 of chapter 4. In the latter firms have an incentive to compete on add-on fees for certain values of l_2, resulting in a more complicated relationship between profits and interfirm distance.

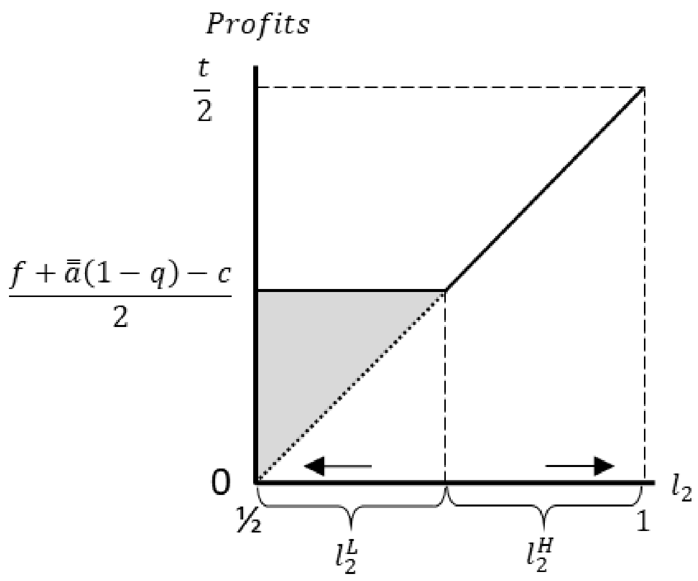

Figure 3.1: Profit and marginal relocation incentives of firm 2 for symmetric locations with $t > \underline{f} + \bar{\bar{a}}(1-q) - c$.

Profits and marginal relocation incentives of firm 1 follow by symmetry. The solid line represents profits of firm 2 for symmetric locations with given l_2. The dotted line represents profits of firm 2 at the sophisticated-consumers benchmark. The vertical distance between the solid and the dotted line represents the additional surplus that firm 2 enjoys due to retained add-on revenues. Given endogenous locations, arrows depict marginal relocation incentives of firm 2 for each given symmetric location pair. An inward-pointing arrow means that firm 2 prefers to offer more a homogeneous product; an outward-pointing arrow means that firm 2 prefers to further differentiate its product. Marginal relocation incentives converge towards the equilibria described in Proposition 2.

While sophisticates always benefit from the presence of naïve types, the extent of the cross-subsidisation they enjoy depends on interfirm distance. To see how, we define the benefit that sophisticated types get from the cross-subsidisation as $b_s = p_s^{bench} - p_s^{eq}$, where p_s^{bench} represents the total price that the average sophisticated consumer would pay at the sophisticated-consumers benchmark, and $p_s^{eq} \in \{p_s^L, p_s^H\}$, with $p_s^{eq} = p_s^L$ if $l_2 \in l_2^L$ and $p_s^{eq} = p_s^H$ if $l_2 \in l_2^H$. Substituting p_s^{bench} and p_s^{eq} and simplifying yields $b_s = \bar{\bar{a}}(1-q)$ if $l_2 \in l_2^H$ and $b_s = c - \underline{f} + t(2l_2 - 1)$ if $l_2 \in l_2^L$. Comparing the two shows that $c - \underline{f} + t(2l_2 - 1)$ lies below $\bar{\bar{a}}(1-q)$ for all $l_2 \in l_2^L$, indicating that sophisticated types enjoy a larger proportion of the naïve consumers' add-on expense when product differentiation is relatively large. This is intuitive, as part of the add-on expense is appropriated by the firms due to the binding price floor constraint when $l_2 \in l_2^L$.

It is worth noting that reducing the proportion of naïve types in the market does not necessarily make sophisticated types worse off. If interfirm distance is such that the price floor constraint is binding ($l_2 \in l_2^L$), marginally increasing q has no effect on the equilibrium upfront price. As such, it has no effect on the extent of the cross-subsidisation that sophisticated types enjoy. However, higher q reduces the parameter range for which the price floor is binding in the first place, as l_2^L falls with q. For $l_2 \in l_2^H$ cross-subsidisation falls with q, as firms do not retain any extra revenue from the naïve types' add-on expense.

Figure 3.2 depicts how the benefit to the average sophisticated consumer from the presence of naïve types depends on interfirm distance. The difference between p_s^{bench} and p_s^{eq} represents the extent of the cross-subsidisation. The vertical distance between p_s^{eq} and the dotted line represents the additional cross-subsidisation that the average sophisticated consumer would enjoy if the price floor constraint was slack. A higher proportion of sophisticated types has two effects: *(i)* it reduces the range l_2^L, as it reduces the parameter space for which the price floor is binding, and *(ii)* shifts the upward-sloping part of p_s^{eq} upwards. The second effect indicates that cross-subsidisation falls with higher q when the price floor constraint is slack.

Note that, despite the lower cross-subsidisation, the average sophisticated consumer pays less at a market with $l_2 \in l_2^L$, as the more intense competition results in lower upfront prices. The relationship between interfirm distance and the surplus of sophisticated (and naïve) consumers is studied in detail in section 3.4.

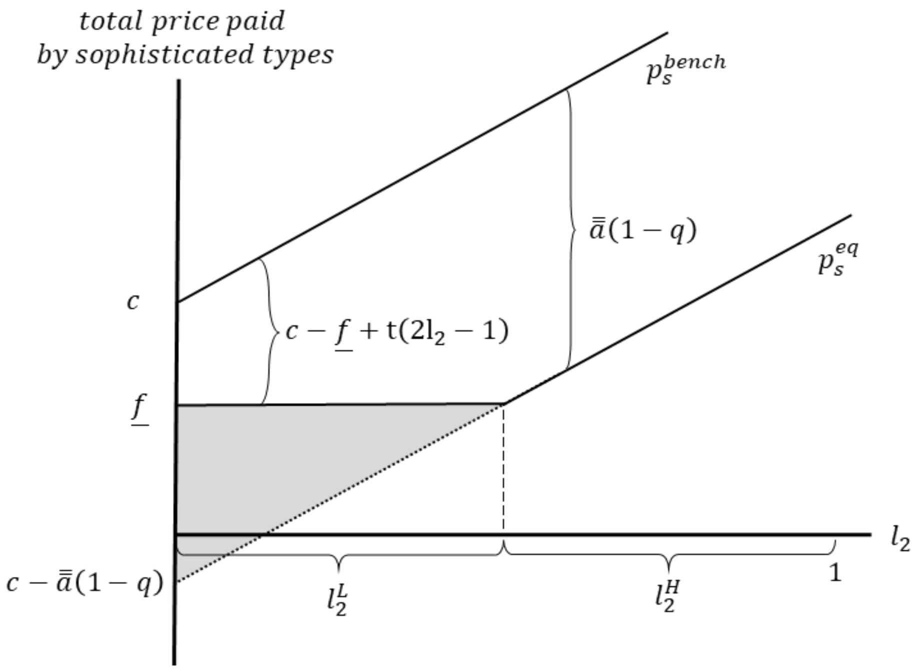

Figure 3.2: Cross-subsidisation of prices paid by the average sophisticated consumer for given symmetric locations.

The vertical distance between p_s^{bench} and p_s^{eq} indicates the benefit that the average sophisticated consumer receives from the cross-subsidisation for any given l_2. The vertical distance between p_s^{eq} and the dotted line represents the additional cross-subsidisation that the average sophisticated consumer would enjoy if the price floor constraint was slack. A higher proportion of sophisticated types (q) has two effects: *(i)* it reduces the range l_2^L, as it reduces the parameter space for which the price floor is binding, and *(ii)* shifts the upward-sloping part of p_s^{eq} upwards. The second effect indicates that benefit to the average sophisticated consumer falls with higher q when the price floor constraint is slack.

3.3.2 Endogenous locations

Corollary 1 established that firms prefer high over low product differentiation. Going further, it implies that firms' profits are the highest with maximum product differentiation. This, however, does not necessarily suggest that firms choose maximum product differentiation in equilibrium when locations are endogenous, as unilateral deviations may be profitable.

Each firm's unilateral choice of whether to deviate from a given location primarily depends on two incentives. The first is the incentive to soften competition by offering a differentiated product – the price effect. The second is the incentive to capture a larger market share by offering a product that is a close substitute to the rival's – the demand effect. The relative strength of these effects determines whether, for given locations, firms prefer to relocate slightly closer to or slightly further apart from their rivals. A well-known result in the standard Hotelling model with variable prices and quadratic distance costs (d'Aspremont *et al.*, 1979; Neven, 1985) is that the price effect dominates the demand effect for any given location pair, establishing a unique subgame perfect equilibrium with maximum differentiation. This is not always the case in our model, as a price effect may or may not be present depending on whether the price floor constraint binds.

The marginal relocation incentives that arise from the interplay between the price and the demand effect are not sufficient to determine the location pairs in equilibrium, as they do not take into account the profitability of large deviations. For instance, if firm 2 is located at, say, $l_2 = 1$, it may not have an incentive to deviate to a location close to 1, but it may have an incentive to deviate to a location next to its rival.

Taking into account both local and large deviation incentives yields two subgame perfect equilibria. They are presented in Proposition 2. Notation *max* and *min* indicates the equilibrium that each variable characterises.

Proposition 2. (Subgame Perfect Nash Equilibria). Let $l_1 = 1 - l_2$.

(a) *If* $\frac{t}{2} \geq \underline{f} + \bar{\bar{a}}(1-q) - c$ *there exists a symmetric maximum-differentiation equilibrium with* $(l_1^{max}, l_2^{max}) = (0,1), (f_{max}, a_{max}) = (c + t - \bar{\bar{a}}(1-q), \bar{\bar{a}})$ *and* $\Pi_{max} = \frac{t}{2}$.

(b) *If* $\underline{f} + \bar{\bar{a}}(1-q) - c \geq \frac{25t}{72}$ *there exists a symmetric minimum-differentiation equilibrium with* $(l_1^{min}, l_2^{min}) = \left(\frac{1}{2}, \frac{1}{2}\right), (f_{min}, a_{min}) = \left(\underline{f}, \bar{\bar{a}}\right)$ *and* $\Pi_{min} = \frac{\underline{f} + \bar{\bar{a}}(1-q) - c}{2}$.

(c) *No other symmetric subgame perfect equilibrium exists.*

Proof: *In appendix B.*

Proposition 2 reveals that in a market in which sophisticated types can avoid add-on fees costlessly, consumer naïvete results in two distinct equilibria – one in which firms offer maximally differentiated products, and one in which firms offer products that are perfect substitutes. To get an intuition behind our model's divergence from the classic result of d'Aspremont *et al.* (1979) and Neven (1985), notice that in the traditional Hotelling model with endogenous prices and quadratic distance costs, the relative strength of the price and the demand effect changes smoothly as firms locate further apart, since profit functions are continuous and differentiable with respect to locations for any (l_i, l_j). The price effect dominates regardless of interfirm distance, establishing a unique maximum-differentiation equilibrium. In contrast, profit functions in our model are kinked, as equilibrium price structure changes depending on interfirm distance. As a result, marginal relocation incentives can differ drastically depending on firm locations.

To see why marginal relocation incentives can point towards location pair $(0,1)$, recall from Proposition 1 that equilibrium upfront prices are increasing in interfirm distance for symmetric location pairs with relatively large interfirm distance ($l_2 \in l_2^H$). Thus, a price effect is present. In particular, the price effect dominates the demand effect, so firms prefer to locate further apart. Intuitively, for $l_2 \in l_2^H$ firms behave as in a standard variable-prices Hotelling model with quadratic distance costs and marginal costs equal to $c - \bar{\bar{a}}(1-q)$. Firms' incentive to locate further apart is exhausted when interfirm distance is maximised – at location pair $(0,1)$.

For an intuition behind why firms may prefer to locate at location pair $\left(\frac{1}{2}, \frac{1}{2}\right)$, recall from Proposition 1 that symmetric location pairs with relatively small interfirm distance ($l_2 \in l_2^L$) result in equilibrium upfront prices equal to the price floor. Since upfront prices are fixed, a price effect does not arise. As a result, each firm's marginal relocation incentive is dominated by the demand effect as in the classic fixed-price Hotelling model. Thus for $l_2 \in l_2^L$ each firm prefers to increase its turf by locating closer to its rival. Firms' incentive to locate closer together is exhausted when interfirm distance is minimised – for location pair $\left(\frac{1}{2}, \frac{1}{2}\right)$.

Figure 3.1, introduced to illustrate how equilibrium profits of firm 2 change with interfirm distance, can be used to depict marginal relocation incentives. For $l_2 \in l_2^L$, firm 2 prefers to locate closer to its rival; this is indicated by an inward-pointing arrow. For $l_2 \in l_2^H$, it prefers to differentiate its product further; this is indicated by an outward-pointing arrow. Marginal relocation incentives of firm 1 follow by symmetry.

Marginal relocation incentives point towards the symmetric location pairs from which firms have no incentive to deviate locally. However, they are not informative regarding the profitability of large deviations. The condition in part *(a)* of Proposition 2 guarantees that a deviation in which firm i locates close to its rival and earns on average $\underline{f} + \bar{\bar{a}}(1-q) - c$ from each consumer in the market yields profits that are lower than Π_{max}. Other large deviations yield even lower profits. This establishes a maximum-differentiation equilibrium. Similarly, condition in part *(b)* of Proposition 2 ensures that a firm which deviates to any location away from the centre of the line in order to enjoy less intense competition cannot earn more than Π_{min}. This establishes a minimum-differentiation equilibrium.

Shifts in the models' parameters can be crucial to the establishment or disestablishment of a certain subgame perfect Nash equilibrium, as they may affect whether conditions *(a)* or *(b)* of Proposition 2 are violated. Higher values of the price floor ($\underline{f}$) and the maximum add-on fee ($\bar{\bar{a}}$) increase the parameter space for which condition *(b)* is satisfied, so they facilitate the existence of the minimum-differentiation equilibrium. Intuitively, this is because they increase the add-on revenues that firms retain when the price floor constraint is binding. The same is true for lower marginal costs (c) and a lower proportion of sophisticated types (q). Higher inherent product differentiability (t) opposes the establishment of a minimum-differentiation equilibrium, as it increases the extent to which a firm can make use of higher product differentiation in order to soften competition. In that way, higher t increases the profits that a firm enjoys by deviating to its end of the line.

Parameter shifts facilitate or inhibit the existence of a maximum-differentiation equilibrium in a similar way. Higher t increases equilibrium profits, as it permits firms to earn more from the maximum interfirm distance. As such, it increases the parameter range for which a maximum-differentiation equilibrium exists – the parameter range for which condition *(b)* in satisfied. Higher $\underline{f}$, $\bar{\bar{a}}$, and lower c, q oppose the establishment of such an equilibrium. This is because they work to increase the profits that a firm can earn by deviating to a location next to its rival.

It is worth noting that the two equilibria coexist for $\frac{t}{2} \geq \underline{f} + \bar{\bar{a}}(1-q) - c \geq \frac{25}{72}t$. This raises the issue of equilibrium selection. We have already established, through Corollary 1, that firm profits are higher when interfirm distance is large. Then it follows that the maximum-differentiation equilibrium is preferred by firms in terms of Pareto dominance when both equilibria exist. A minimum-differentiation equilibrium is inferior despite firms earning more than in the sophisticated-consumers benchmark. However, it is not obvious that firms would find it easy to coordinate towards a maximum-differentiation equilibrium. The prospect of such coordination appears less likely when one considers possible costs associated with differentiating one's product (Grossman and Shapiro, 1984).

Going beyond the issue of equilibrium selection when the two multiple equilibria arise, it is worth highlighting the market parameters that establish the most favourable outcome for firms. We focus our attention on the two parameters that are the central to the academic and legislative discussion on markets with add-on fees: the size of the add-on fees ($\bar{\bar{a}}$) and the proportion of sophisticated consumers in the market (q).[39] Taken together, these parameters characterise the average add-on revenue per consumer that each firm earns, $\bar{\bar{a}}(1-q)$. Conventional wisdom suggests that firms prefer higher over lower $\bar{\bar{a}}(1-q)$. Figure 3.3 shows that this is not always the case.

Figure 3.3 plots profits at the maximum-differentiation and the minimum-differentiation equilibrium for different values of $\bar{\bar{a}}(1-q)$. The dashed line represents Π_{min} for the values of $\bar{\bar{a}}(1-q)$ for which the minimum-differentiation equilibrium exists. The solid line represents Π_{max}. Note that Π_{min} is increasing in $\bar{\bar{a}}(1-q)$, as a higher $\bar{\bar{a}}(1-q)$ allows firms to enjoy higher add-on revenues without subsidising lower upfront prices. Π_{max} is independent of $\bar{\bar{a}}(1-q)$, as subsidisation takes place. However, Π_{max} is greater than Π_{min} for certain parameter values. Hence, whether firms prefer higher $\bar{\bar{a}}(1-q)$ depends on whether the higher $\bar{\bar{a}}(1-q)$ establishes a minimum-differentiation equilibrium, and whether that equilibrium features profits above Π_{max}. Values of $\bar{\bar{a}}(1-q)$ such that $\bar{\bar{a}}(1-q) \geq c + t - \underline{f}$ establish the best outcome for firms as this condition guarantees the existence of a minimum-

[39] We are discussing the desirability of market policies that affect these two parameters in detail in the following section.

differentiation equilibrium with $\Pi_{min} > \Pi_{max}$. When that condition takes places, firms strictly prefer to earn higher average add-on revenue per consumer. If condition $\bar{\bar{a}}(1-q) \geq c + t - \underline{f}$ cannot be satisfied (perhaps due to existing legislation) then firms are better off with $\bar{\bar{a}}(1-q)$ low enough to establish a maximum-differentiation equilibrium, as they can benefit more from the associated higher upfront prices.

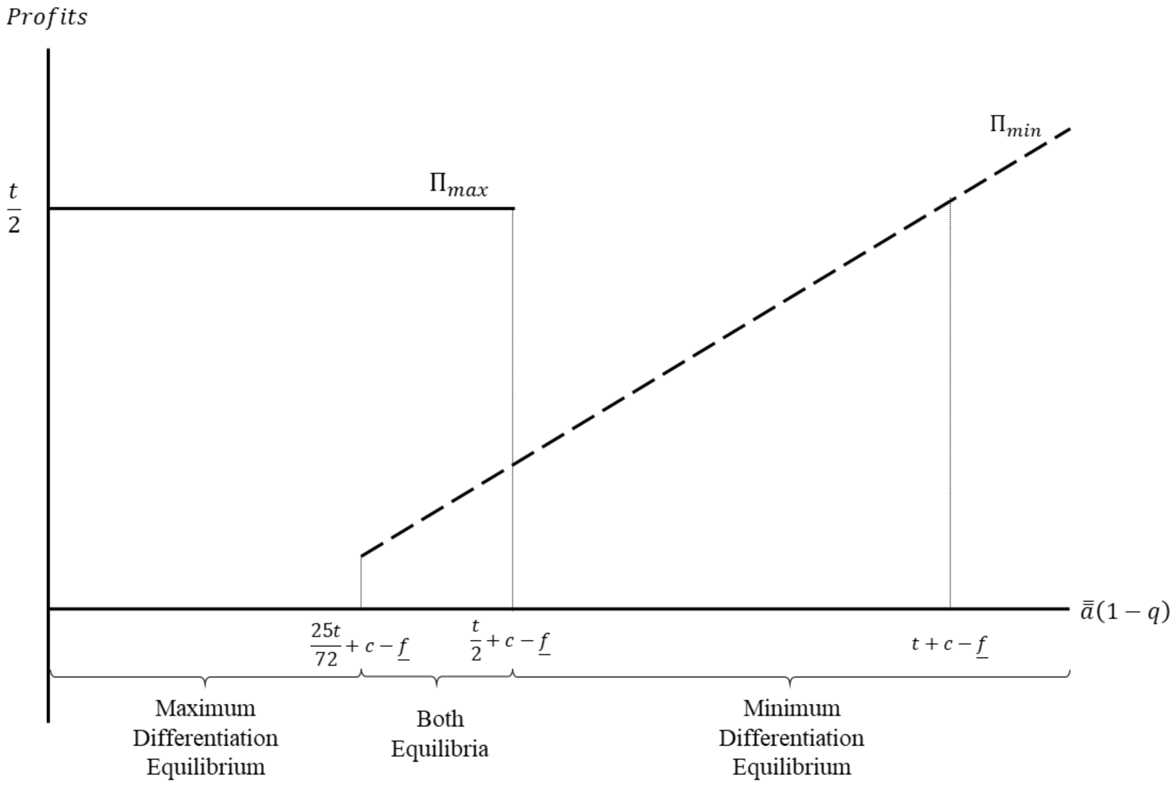

Figure 3.3: Firm profits for given add-on revenues.

The solid line represents profits at the maximum-differentiation equilibrium (Π_{max}), for all values of $\bar{\bar{a}}(1-q)$ for which this equilibrium exists. The dashed line represents profits at the minimum-differentiation equilibrium (Π_{min}), for all values of $\bar{\bar{a}}(1-q)$ for which this equilibrium exists. Π_{min} is increasing in $\bar{\bar{a}}(1-q)$, as a higher $\bar{\bar{a}}(1-q)$ allows firms to enjoy higher add-on revenues without subsidising lower upfront prices. Π_{max} is independent of $\bar{\bar{a}}(1-q)$, as subsidisation takes place. However, Π_{max} is greater than Π_{min} for any $\bar{\bar{a}}(1-q)$ below $t + c - \underline{f}$.

3.4 Welfare and policy implications

Markets in which consumers underestimate future costs often face regulatory scrutiny, as firms have an incentive to exploit consumer biases. The 2009 Credit Card Accountability Responsibility and Disclosure (CARD) Act in the United States is recent example of a regulation aimed at protecting consumers. Among other interventions, it permits firms to charge overlimit fees only if consumers explicitly opt in to such service, and introduces a $25 limit on late fees. In our framework, the first intervention would describe an increase in the proportion of sophisticated types, q. The second intervention would be interpreted as a reduction in the maximum fees that consumers may pay in any scenario, $\bar{\bar{a}}$. In this section we explore the extent to which similar policies can be used to improve market outcomes in the context of our model.[40]

In particular, we study the effect of similar policies on the welfare of the average consumer, as well as on the welfare of naïve and sophisticated types separately. With consumers distributed on a Hotelling line and full market coverage, consumer welfare depends on the prices that the average consumer pays, as well as on the travel costs she incurs to purchase from her firm of choice. However, average travel costs are always maximised in our model, as the two equilibria are characterised by either maximum or minimum product differentiation. As such we ignore this dimension when studying changes in consumer surplus. The fact that average travel costs are maximised in every equilibrium also implies that the two equilibria are identical in terms of total surplus, and that conventional regulatory policies cannot improve market outcomes with respect to that standard. Hence, studying how regulatory policies can affect total surplus lies outside the scope of our work.[41]

Since average travel costs do not play a role in our model, we study consumer welfare by focusing on the effect of market policies on the total price paid by the average consumer, the total price paid by each sophisticated type, and the total price paid by each naïve type. For a market at the maximum-differentiation equilibrium,

[40] See Fletcher (2016) for a review on the effectiveness of various types of market remedies.

[41] Analysis of how market remedies can affect average travel costs takes place in Chapter 4, as the model studied there features an equilibrium with intermediate product differentiation.

average total price is given by $p^{max} = f_{max} + a_{max}(1-q) = c + t$. Similarly, for a market at the minimum-differentiation equilibrium, average total price is given by $p^{min} = f_{min} + a_{min}(1-q) = \underline{f} + \bar{\bar{a}}(1-q)$. Sophisticated consumers pay only the upfront price. As such, the total price paid by the average sophisticated consumer is given by $p_s^{max} = f_{max} = c + t - \bar{\bar{a}}(1-q)$ at the maximum-differentiation equilibrium and by $p_s^{min} = f_{min} = \underline{f}$ at the minimum-differentiation equilibrium. Naïve types pay both price components. Hence, the total price paid by the average naïve consumer is given by $p_n^{max} = f_{max} + a_{max} = c + t + \bar{\bar{a}}q$ at the maximum-differentiation equilibrium, and by $p_n^{min} = f_{min} + a_{min} = \underline{f} + \bar{\bar{a}}$ at the minimum-differentiation equilibrium.

We investigate the effectiveness of two well-known categories of policies. The first category is disclosure policies that aim to educate consumers about the relevance of add-on fees. In our model, this results in a higher proportion of sophisticated types (q) in the market. For example, a market remedy that requires banks to send text alerts to consumers that are about to go into unarranged overdraft falls in this category. The second category is price control policies – policies that impose a cap on the add-on fees that naïve consumers pay. In our model this is translated as a reduction in $\bar{\bar{a}}$. Such a policy would, for example, place a limit on the fees that consumers can accumulate by being late on their credit card debt payments.

In order to explore the effectiveness of each policy in detail, we separately study how that policy affects the average consumers, naïve consumers, and sophisticated consumers. When discussing disclosure policies, we borrow the terminology of Gabaix and Laibson (2006) and refer to naïve consumers that become sophisticated after the policy takes place as "informed naïve consumers". Accordingly, consumers that remain naïve are referred to as "uninformed naïve consumers". We refer to consumers that were sophisticated before the policy takes place simply as "sophisticated consumers".

3.4.1 Policies that maintain the market equilibrium

We first discuss the desirability of policies that do not introduce a shift in the market equilibrium. As a first order effect, a disclosure policy (higher q) reduces the proportion of consumers that pay add-on fees in equilibrium. This strictly improves the welfare of the informed naïve consumers in both equilibria, since $p_n^{max} > p_s^{max}$ and $p_n^{min} > p_s^{min}$.

The same does not necessarily apply to sophisticated consumers or uninformed naïve consumers. For a market in a maximum-differentiation equilibrium, both sophisticated consumers and uninformed naïve consumers become worse off with higher q, as the resulting lower add-on revenues reduce the subsidisation on upfront prices. To see this notice that p_n^{max} and p_s^{max} are increasing in q. For a market in the minimum-differentiation equilibrium, sophisticated consumers and uninformed naïve consumers are unaffected by a disclosure policy (p_n^{min} and p_s^{min} are independent of q). This is because the subsidisation of upfront prices is exhausted by the binding price floor, so any reduction in add-on revenues only results in lower firm profits.

The impact of a disclosure policy on the welfare of the average consumer depends crucially on whether firms have an incentive respond to the policy by passing the reduction of add-on revenues back to consumers through setting higher upfront prices. Such supply-side response takes place in the maximum-differentiation equilibrium, as equilibrium upfront prices lie above the price floor. As a result, a higher q does not affect the average total price (p^{max} is independent of q). Thus, any benefit to the educated consumers is funded by the higher prices that are paid by other consumers. In the minimum-differentiation equilibrium, firms have no incentive to respond to the policy by increasing upfront prices, as upfront prices are artificially high due to the binding price floor constraint. As a result, increasing q lowers the add-on revenues that firms retain in equilibrium. Thus, any benefit to the educated consumers is funded by lowering firms' profits.

A policy that lowers $\bar{\bar{a}}$ has the direct effect of lowering the add-on fees that naïve consumers pay in both equilibria. As a result, naïve consumers benefit by paying a lower total price (p_n^{max}, p_n^{min} are increasing in $\bar{\bar{a}}$). However, the benefit they enjoy is not the same for each equilibrium; it depends on whether firms respond to the lower add-on fees by reducing the subsidy on upfront prices. This response takes place at the maximum-differentiation equilibrium, as the price floor constraint is not binding. The resulting decrease in the subsidy reduces the benefit to naïve consumers. However the benefit is not fully negated, as the loss due to the higher upfront prices is shared evenly across consumers. As a result, lowering $\bar{\bar{a}}$ decreases p_n^{max} by a factor q. In contrast, naïve consumers fully enjoy the benefit of a lower $\bar{\bar{a}}$ in a minimum-differentiation equilibrium, as firms have no incentive to respond to the policy by increasing upfront prices above the binding price floor. As such, p_n^{min} decreases by a factor 1.

A reduction in $\bar{\bar{a}}$ affects total prices paid by sophisticated consumers only through the subsidy on upfront prices. In the maximum-differentiation equilibrium a lower $\bar{\bar{a}}$ increases the price sophisticated types are paying (p_s^{max} is decreasing in $\bar{\bar{a}}$), as it reduces the subsidy on upfront prices. In the minimum-differentiation equilibrium firms have no incentive to increase prices above the price floor. As such, decreasing $\bar{\bar{a}}$ has no effect on the price paid by sophisticates (p_s^{min} in independent of $\bar{\bar{a}}$).

The impact of a decrease in $\bar{\bar{a}}$ on the welfare of the average consumer again depends crucially on whether firms increase upfront prices in response. In the maximum-differentiation equilibrium the total price paid by the average consumer does not change, as firms increase upfront prices enough to compensate. Thus, the benefit the naïve types receive by the lower add-on fees is funded by the higher upfront prices that each consumer pays. In the minimum-differentiation equilibrium firms have no incentive to respond by increasing upfront prices. As a result, lowering $\bar{\bar{a}}$ results in lower retained add-on revenues. Thus, the benefit to the naïve types is funded by the decrease in firm profits.

3.4.2 Policies that shift the market equilibrium

We now consider policies extensive enough to shift the market equilibrium. As with policies that marginally affect market parameter, a policy that shifts a market equilibrium may benefit some consumers to the detriment of others. Proposition 3 collects our findings regarding such equilibrium-shifting market policies.

Proposition 3. (Price controls and disclosure policies)

(i) *For a market at a maximum-differentiation equilibrium, policies that decrease* $\bar{\bar{a}}$ *and increase* q *have only marginal effects.*

(ii) *Consider a market at a minimum-differentiation equilibrium with* q_{old} *sophisticated types, and a disclosure policy that results in* q_{new} *sophisticated types and establishes a maximum-differentiation equilibrium, where* $q_{new} > q_{old}$*. The disclosure policy:*

(a) Benefits the average consumer if $\underline{f} + \bar{\bar{a}}(1 - q_{old}) > c + t$

(b) Harms sophisticated consumers

(c) Harms uninformed naïve consumers

(d) Benefits informed naïve consumers if q_{new} *satisfies* $\underline{f} + \bar{\bar{a}} > c + t - \bar{\bar{a}}(1 - q_{new})$

(iii) *Consider a market at a minimum-differentiation equilibrium with add-on fees* $\bar{\bar{a}}_{old}$*, and a price control policy that results in add-on fees equal to* $\bar{\bar{a}}_{new}$ *and establishes a maximum-differentiation equilibrium, where* $\bar{\bar{a}}_{new} < \bar{\bar{a}}_{old}$*. The price control policy:*

(a) Benefits the average consumer if $\underline{f} + \bar{\bar{a}}_{old}(1 - q) > c + t$

(b) Harms sophisticated consumers

(c) Benefits naïve consumers if $\bar{\bar{a}}_{old}$ *and* $\bar{\bar{a}}_{new}$ *satisfy* $\underline{f} + \bar{\bar{a}}_{old} > c + t + \bar{\bar{a}}_{new}q$

Proof: *In appendix B.*

For an intuition behind part *(i)* of Proposition 3, notice that conventional disclosure and price control policies can only lead to a new market equilibrium if the market is originally at a minimum-differentiation equilibrium. This is because, as we showed in Proposition 2, lower $\bar{\bar{a}}$ and higher q both facilitate the existence of a maximum-differentiation equilibrium and inhibit the existence of a minimum-differentiation equilibrium.

3.4.2.1 *Effect on the average consumer*

Given that higher equilibrium profits only arise from a higher average total price, the question whether a shift to the maximum-differentiation equilibrium benefits the average consumer is answered in Figure 3.3. A shift to the maximum-differentiation equilibrium leads to lower firm profits only if the pre-intervention market parameters satisfy $\bar{\bar{a}}(1-q) \geq c+t-\underline{f}$, as in that case, total average prices at a minimum-differentiation equilibrium are higher than prices with maximum product differentiation. Conditions in parts *(ii-a)* and *(iii-a)* rearrange condition $\bar{\bar{a}}(1-q) \geq c+t-\underline{f}$ to more clearly compare the total average pre-intervention price (given minimum product differentiation) with the total average post-intervention price (given maximum product differentiation).

3.4.2.2 *Effect on naïve and sophisticated consumers*

A policy that shifts the market equilibrium does not affect naïve and sophisticated types in the same way. Furthermore, whether each types is made better off depends on whether the shift to a maximum-differentiation equilibrium is the result of a disclosure or a price control policy. Since sophisticated consumers avoid any add-on fees, disclosure policies and price controls provide no benefit to them. In fact, sophisticated consumers are harmed by a shift to the maximum-differentiation equilibrium, as that equilibrium features higher upfront prices than the minimum-differentiation equilibrium. A disclosure policy that establishes a maximum-differentiation equilibrium makes uninformed naïve consumers worse off, as it results in higher upfront prices, while leaving add-on prices unaffected. Whether informed naïve consumers become better off depends on whether the benefit from avoiding add-on fees outweighs the loss from paying higher upfront prices. Intuitively, the higher the q_{new}, the lower is the benefit to the informed naïves. This is because a smaller number of uninformed naïve types in the market results in lower add-on revenues, and by extension, smaller subsidies on upfront prices. A price control policy that establishes a maximum-differentiation equilibrium improves the position of the average naïve consumer if the reduction in the add-on fees that she pays can make up for the higher upfront prices associated with the maximum-differentiation equilibrium.

Taken together with the findings of section 3.4.1, Proposition 3 shows that a change in product locations as a supply-side response can drastically affect the effectiveness of a policy intervention, as it can undo any benefits to consumers or even result in further harm. In particular, a regulator that takes such a supply-side response into account may find that policies that have a large effect on market parameters may be less beneficial than policies that have a small effect on market parameters. For example, given a market at the minimum-differentiation equilibrium, a decrease in add-on fees benefits naïve consumers and has no effect on the prices paid by sophisticated consumers if the decrease is small enough that the market remains at the minimum-differentiation equilibrium. However, decreasing add-on fees enough to shift the market to the maximum-differentiation equilibrium harms sophisticated consumers and may even harm naïve consumers, due to the associated higher upfront prices.

Overall, in agreement with past literature, Proposition 3 shows that standard regulatory policies will create winners and losers. Thus, a regulator needs to look beyond the impact on the welfare of the average consumer and pay attention to how a policy distributes surplus between naïve and sophisticated types.

Going beyond the different impact on naïve and sophisticated types, designing policies solely aimed to benefit the average consumer is not without issues. Parts *(ii-a)* and *(iii-a)* of Proposition 3 state that interventions which shift the market to a maximum-differentiation equilibrium may be beneficial. The maximum-differentiation equilibrium is, however, not the best outcome for the average consumer. As long as post-intervention parameters satisfy $\bar{\bar{a}}(1-q) < c + t - \underline{f}$, a market policy that results in a minimum-differentiation equilibrium makes her better off. This is because the minimum-differentiation equilibrium with such parameters features lower total average prices compared to the maximum-differentiation equilibrium. To see this, notice that Π_{min} is below Π_{max} in Figure 3.3 for all $\bar{\bar{a}}(1-q) < c + t - \underline{f}$.

Note that shifting the market from a maximum-differentiation to a minimum-differentiation equilibrium requires a rather unconventional policy. This is because a minimum-differentiation is established when the proportion of sophisticated consumers (q) is relatively small, and the add-on fees ($\bar{\bar{a}}$) are relatively high. Such an unconventional policy could, for instance, permit firms to present various contingency charges using smaller fonts, reducing the attention that consumers may pay to them. It could also increase the cap on such charges, permitting firms to earn more from naïve types.

Finally, viewing Figure 3.3 from the perspective of the average consumer also reveals that standard measures of market performance, such as supra-competitive profits, may not be revealing of the markets in which the average consumer is harmed the most. A seemingly competitive market – one in which products are perfect substitutes – may be considered problematic if firms earn supra-competitive profits by retaining add-on revenue. However, this is the best outcome for consumers if market parameters satisfy $\bar{\bar{a}}(1-q) < c + t - \underline{f}$. In contrast, a market in which firms offer maximally differentiated products and earn profits that are analogous to the degree of product differentiation may be considered to be performing well, despite consumers getting a worse deal compared to a market with perfect substitutes and supra-competitive profits.

3.5 Conclusion

When consumers make mistakes, firms may adjust choice variables in order to exploit them. This chapter studies how the possibility of profitably exploiting consumer mistakes impacts such a choice variable – horizontal product differentiation.

We show that the degree of product differentiation critically influences price structure and the size of extra profits that firms can retain from naïve consumers that accidentally buy unwanted add-ons. High differentiation yields high prices, but results in full dissipation of add-on revenues, while low differentiation permits retention of some extra profits, but results in lower upfront prices. This more complex relationship between product differentiation and profitability can incentivise firms to offer less differentiated products and may, under certain conditions, establish a minimum-differentiation equilibrium. While firms at the minimum-differentiation equilibrium profit from the naïve consumers’ add-on expenditure, they may earn less compared to a market populated solely by sophisticated types.

We show that exploiting naïve consumer may have a silver lining, as the resulting competition on product locations may outweigh consumer harm from unanticipated prices. Market policies can harm consumers if they weaken firms' incentives to compete on product locations. The overall effectiveness of interventions depends on the pre-intervention market equilibrium, as well as on whether interventions can shift the market to a different equilibrium.

Being one of the first to address the importance of product differentiation when analysing markets with naïve consumers, our model leaves many questions unanswered. For example, in considering an exogenous proportion of naïve consumers it does not explore the incentives of firms to educate or further confuse consumers. It also does not take into account the various consumption distortions that may arise from below-cost upfront prices and above-cost add-on fees. We leave these questions for future work.

3.6 Appendix B

Proof of Proposition 1.

Given symmetric locations ($l_1 = 1 - l_2$), best response functions become

$$f_1^{BR} = \max\left(\frac{t(2l_2 - 1) + c - \bar{\bar{a}}(1-q) + f_2}{2}, \underline{f}\right)$$

$$f_2^{BR} = \max\left(\frac{t(2l_2 - 1) + c - \bar{\bar{a}}(1-q) + f_1}{2}, \underline{f}\right)$$

Symmetric locations imply that both firms play the same equilibrium upfront price, f^*. Suppose that the price floor is not binding. Then f^* satisfies $f^* = \frac{t(2l_2-1)+c-\bar{\bar{a}}(1-q)+f^*}{2}$. Solving gives $f^* = t(2l_2 - 1) + c - \bar{\bar{a}}(1-q)$. The price floor constraint is not binding as long as $f^* > \underline{f}$. Solving the inequality for l_2 gives that the price floor is not binding if $l_2 \in \left[\frac{1}{2} + \frac{\underline{f}+\bar{\bar{a}}(1-q)-c}{2t}, 1\right] = l_2^H$. Thus, for $l_2 \in l_2^H$ there exists a pure strategy price equilibrium with $f^* = t(2l_2 - 1) + c - \bar{\bar{a}}(1-q) = f_H^*$. The price floor is binding if $l_2 \in \left[\frac{1}{2}, \frac{1}{2} + \frac{\underline{f}+\bar{\bar{a}}(1-q)-c}{2t}\right) = l_2^L$. Thus for $l_2 \in l_2^L$ there exists a pure strategy price equilibrium with $f^* = \underline{f} = f_L^*$. To find equilibrium profits Π_H^*, Π_L^* replace f_H^* and f_L^* in place of f_1 and f_2 in $\Pi_1 = (f_1 + \bar{\bar{a}}(1-q) - c)x$ and $\Pi_2 = (f_2 + \bar{\bar{a}}(1-q) - c)(1-x)$.■

Proof of Corollary 1.

For part *(a)*, see that $\frac{d\Pi_H^*}{dl_2} > 0$, while $\frac{d\Pi_L^*}{dl_2} = 0$. For part *(b)*, recall that symmetric location pairs with $l_2 \in l_2^H$ yield equilibrium profits Π_H^*, while symmetric location pairs with $l_2 \in l_2^L$ yield equilibrium profits Π_L^*. Comparison between the two yields that $\Pi_H^* > \Pi_L^*$. Since, by part *(a)*, profits in $l_2 \in l_2^L$ are independent of l_2, the statement of part *(b)* is true. ■

Proof of Proposition 2.

The proof proceeds in 4 steps. In step 1 we solve the price equilibrium for general locations. The price equilibrium is unique and comes in four types, depending on firm locations. In step 2 we examine marginal relocation incentives from any given symmetric location pair. Marginal relocation incentives converge towards location pairs $(0,1)$ and $\left(\frac{1}{2},\frac{1}{2}\right)$. In step 3 we explore the profitability of large deviations from pair $(0,1)$ and arrive at the condition in part *(a)* of the Proposition. In step 3 we explore the profitability of large deviations from pair $\left(\frac{1}{2},\frac{1}{2}\right)$ and arrive at the condition in part *(b)* of the Proposition.

Step 1. *Price equilibrium for general locations.*

To characterise the equilibrium locations we first solve the pricing stage for general locations. Suppose firm locations are such that the price floor constraint is not binding for either firm; that is $f_1^{BR}=\frac{t(l_1+l_2)(l_2-l_1)+c-\bar{\bar{a}}(1-q)+f_2}{2}$ and $f_2^{BR}=\frac{t(2-l_1-l_2)(l_2-l_1)+c-\bar{\bar{a}}(1-q)+f_1}{2}$. Simultaneously solving for f_1^*,f_2^* yields a unique solution $\left(f_{1\ NB}^*,f_{2\ NB}^*\right)=\left(c-\bar{\bar{a}}(1-q)+\frac{t}{3}(l_2-l_1)(2+l_1+l_2),c-\bar{\bar{a}}(1-q)+\frac{t}{3}(l_2-l_1)(4-l_1-l_2)\right)$, where notation *NB* denotes a non-binding price floor constraint for both firms. For $\left(f_{1\ NB}^*,f_{2\ NB}^*\right)$ to describe an equilibrium it must hold that the price floor constraint does not bind for either firm; that is $\underline{f}\leq c-\bar{\bar{a}}(1-q)+\frac{t}{3}(l_2-l_1)(2+l_1+l_2)$ and $\underline{f}\leq c-\bar{\bar{a}}(1-q)+\frac{t}{3}(l_2-l_1)(4-l_1-l_2)$.

Now suppose that the price floor constraint is binding for firm 1 only; that is $f_1^{BR}=\underline{f}$ and $f_2^{BR}\geq\underline{f}$. Then, simultaneously solving the system of best response functions yields a unique solution $\left(f_{1\ AB1}^*,f_{2\ AB1}^*\right)=\left(\underline{f},c-\bar{\bar{a}}(1-q)+\frac{t}{3}(l_2-l_1)(4-l_1-l_2)\right)$, where notation *AB1* denotes a price floor constraint that is only binding for firm 1. For $\left(f_{1\ AB1}^*,f_{2\ AB1}^*\right)$ to describe an equilibrium it must hold that the price floor constraint binds for firm 1 only; that is $\frac{t(l_1+l_2)(l_2-l_1)+c-\bar{\bar{a}}(1-q)+f_2^*}{2} < \underline{f}$ and $\frac{t(2-l_1-l_2)(l_2-l_1)+c-\bar{\bar{a}}(1-q)+f_1^*}{2} \geq \underline{f}$.

Next suppose that the price floor constraint is binding for firm 2 only; that is $f_2^{BR} = \underline{f}$ and $f_1^{BR} \geq \underline{f}$. Then, simultaneously solving the system of best response functions yields a unique solution $\left(f_1^*{}_{AB2}, f_2^*{}_{AB2}\right) = \left(\frac{t(l_1+l_2)(l_2-l_1)+c-\bar{a}+\underline{f}}{2}, \underline{f}\right)$, where notation *AB2* denotes a price floor constraint that is only binding for firm 2. For $\left(f_1^*{}_{AB2}, f_2^*{}_{AB2}\right)$ to describe an equilibrium it must hold that the price floor constraint binds for firm 2 only; that is $\frac{t(l_1+l_2)(l_2-l_1)+c-\bar{\bar{a}}(1-q)+f_2^*}{2} \geq \underline{f}$ and $\frac{t(2-l_1-l_2)(l_2-l_1)+c-\bar{\bar{a}}(1-q)+f_1^*}{2} < \underline{f}$.

Finally, suppose that the price floor constraint is binding for both firms; that is $f_1^{BR} = f_2^{BR} = \underline{f}$. This implies an equilibrium price pair $\left(f_1^*{}_{BB}, f_2^*{}_{BB}\right) = \left(\underline{f}, \underline{f}\right)$, where notation *BB* denotes a price floor constraint that is binding for both firms. For this to be the case it must hold that $\frac{t(l_1+l_2)(l_2-l_1)+c-\bar{\bar{a}}(1-q)+f_2^*}{2} < \underline{f}$ and $\frac{t(2-l_1-l_2)(l_2-l_1)+c-\bar{\bar{a}}(1-q)+f_1^*}{2} < \underline{f}$.

Step 2. *Marginal relocation incentives converge towards location pairs* $(0,1)$ *and* $\left(\frac{1}{2}, \frac{1}{2}\right)$.

In order to calculate the marginal relocation incentives of each firm for a given symmetric location pair we need to express equilibrium profits in terms of l_1, l_2.

First, consider symmetric location pairs that satisfy $l_2 \in l_2^H$. Applying $l_1 = 1 - l_2$ and $l_2 \in l_2^H$ in the conditions of step 1 gives that for symmetric locations with $l_2 \in l_2^H$, the price equilibrium is of type NB. Substituting $\left(f_1^*{}_{NB}, f_2^*{}_{NB}\right)$ at each firm's profit function gives that firm 1 earns profit $\Pi_{H1}^* = \frac{t(l_2-l_1)(2+l_1+l_2)^2}{18}$, while firm 2 earns profit $\Pi_{H2}^* = \frac{t(l_2-l_1)(4-l_1-l_2)^2}{18}$. Differentiation yields $\frac{d\Pi_{H1}^*}{dl_1} < 0$ and $\frac{d\Pi_{H2}^*}{dl_2} > 0$; that is, each firms prefers to locate further away from its rival. Now notice that (l_1^{max}, l_2^{max}) satisfies $l_2 \in l_2^H$. Since (l_1^{max}, l_2^{max}) describes locations at the ends of the line we can conclude that any local deviation from (l_1^{max}, l_2^{max}) is unprofitable, as firms cannot locate any further apart. Replacing l_1^{max} and l_2^{max} in equilibrium prices and profits gives the desired f_{max}, a_{max}, and Π_{max}.

Next, consider symmetric location pairs that satisfy $l_2 \in l_2^L$. Applying $l_1 = 1 - l_2$ and $l_2 \in l_2^L$ in the conditions of step 1 gives that for symmetric locations with $l_2 \in l_2^L$, the price equilibrium is of type BB. Substituting $\left(f_{1\,BB}^*, f_{2\,BB}^*\right)$ at each firm's profit function gives that firm 1 earns profit $\Pi_{L1}^* = \frac{\left(\underline{f}+\bar{\bar{a}}(1-q)-c\right)(l_1+l_2)}{2}$, while firm 2 earns profit $\Pi_{L2}^* = \frac{\left(\underline{f}+\bar{\bar{a}}(1-q)-c\right)(2-l_1-l_2)}{2}$. Differentiation yields $\frac{d\Pi_{L1}^*}{dl_1} > 0$ and $\frac{d\Pi_{L2}^*}{dl_2} < 0$; that is, each firms prefers to locate closer to its rival. But $\left(l_1^{min}, l_2^{min}\right)$ satisfies $l_2 \in l_2^L$ and describes locations at the centre of the line. This any local deviation from $\left(l_1^{min}, l_2^{min}\right)$ is unprofitable, as firms cannot locate any further apart. Replacing l_1^{min} and l_2^{min} in equilibrium prices and profits gives the desired f_{min}, a_{min}, and Π_{min}.

Step 3. *Profitability of large deviations from location pair* $(0,1)$.

Let $l_2 = 1$ and consider a deviation at some $l_1^{d^{\max}} \in (0,1]$. Deviation incentives of firm 2 follow in an analogous manner by relabelling firms. Deviation profits of firm 1 are given by $\Pi_{1\,NB}^{d^{max}} = \left(f_{1\,NB}^* + a - c\right)\left(\frac{l_1^{d^{\max}}}{2} + \frac{1}{2} - \frac{f_{1\,NB}^* - f_{2\,NB}^*}{2t\left(1-l_1^{d^{\max}}\right)}\right)$ if $l_1^{d^{\max}}$ is such that the price equilibrium is of type *NB*; that is if $l_1^{d^{\max}} \in \left(0, 2 - \sqrt{1 + \frac{3\left(\underline{f}+\bar{\bar{a}}(1-q)-c\right)}{t}}\right] = l_{1\,NB}^{d^{\max}}$. Differentiating with respect to l_1 yields that $\frac{d\Pi_{1\,NB}^{d^{max}}}{dl_1^{d^{\max}}} = \left(f_{1\,NB}^* + \bar{\bar{a}}(1-q) - c\right)\left(\frac{1}{2} - \frac{2}{3\left(1-l_1^{d^{\max}}\right)}\right) < 0$. Thus, deviating to any $l_1^{d^{\max}} \in l_{1\,NB}^{d^{\max}}$ is less profitable than choosing $l_1 = 0$.

Deviation profits are given by $\Pi_{1\,BB}^{d^{max}} = \left(\underline{f} + \bar{\bar{a}}(1-q) - c\right)\left(\frac{l_1^{d^{\max}}}{2} + \frac{1}{2}\right)$ if $l_1^{d^{\max}}$ is such that the price equilibrium is of type *BB*; that is if $l_1^{d^{\max}} \in \left(\sqrt{1 - \frac{\left(\underline{f}+\bar{\bar{a}}(1-q)-c\right)}{t}}, 1\right] = l_{1\,BB}^{d^{\max}}$. Differentiating with respect to l_1 yields that

$\frac{d\Pi_{1\,BB}^{d^{max}}}{dl_1^{d^{\max}}} = \frac{\left(\underline{f}+\bar{\bar{a}}(1-q)-c\right)}{2} > 0$. This implies that among the $l_1^{d^{\max}} \in l_{1\,BB}^{d^{\max}}$, the optimal deviation for firm 1 is at $l_1^{d^{\max}} = 1-\epsilon$, where ϵ is an infinitesimal positive value. At $l_1^{d^{\max}} = 1-\epsilon$, deviation profit is almost equal to $\underline{f}+\bar{\bar{a}}(1-q)-c$. Therefore, this deviation is profitable if $\underline{f}+\bar{\bar{a}}(1-q)-c > \frac{t}{2}$.

Deviation profits are given by $\Pi_{1\,AB2}^{d^{max}} = \left(f_{1\,AB2}^{*}+\bar{\bar{a}}(1-q)-c\right)\left(\frac{l_1^{d^{\max}}}{2}+\frac{1}{2}-\frac{f_{1\,AB2}^{*}-\underline{f}}{2t\left(1-l_1^{d^{\max}}\right)}\right)$ if $l_1^{d^{\max}}$ is such that the price equilibrium is of type *AB2*; that is if $l_1^{d^{\max}} \in \left(2-\sqrt{1+\frac{3\left(\underline{f}+\bar{\bar{a}}(1-q)-c\right)}{t}}, \sqrt{1-\frac{\left(\underline{f}+\bar{\bar{a}}(1-q)-c\right)}{t}}\right] = l_{1\,AB2}^{d^{\max}}$. Differentiating with respect to l_1 yields that $\frac{d\Pi_{1\,AB2}^{d^{max}}}{dl_1^{d^{\max}}} = \left(f_{1\,AB2}^{*}+\bar{\bar{a}}(1-q)-c\right)\left(\frac{1}{2}-\frac{t\left(1-l_1^{d^{\max 2}}\right)-\left(\underline{f}+\bar{\bar{a}}(1-q)-c\right)}{4t\left(1-l_1^{d^{\max}}\right)^2}\right)$. For $t < 3\left(\underline{f}+\bar{\bar{a}}(1-q)-c\right)$, it can be seen that $\frac{d\Pi_{1\,AB2}^{d^{max}}}{dl_1^{d^{\max}}} > 0$. Therefore, deviation profits at $l_1^{d^{\max}} \in l_{1\,AB2}^{d^{\max}}$ are lower than deviation profits at $l_1^{d^{\max}} \in l_{1\,BB}^{d^{\max}}$, which we examined above. For $t \geq 3\left(\underline{f}+\bar{\bar{a}}(1-q)-c\right)$, deviation profits have a local maximum at $l_1^{d^{\max}} = \frac{2}{3}-\frac{\sqrt{1-\frac{3\left(\underline{f}+\bar{\bar{a}}(1-q)-c\right)}{t}}}{3}$ or a corner solution. After some algebra it can be seen that neither the local maximum nor any corner solution constitutes a profitable deviation, as deviation profits are always below Π_{max}.

Finally, there does not exist an l_1^d such that the price equilibrium is of type *AB1*. Therefore, there exists a maximum-differentiation equilibrium if $\frac{t}{2} \geq \underline{f}+\bar{\bar{a}}(1-q)-c$.

Step 4. *Profitability of large deviations from location pair* $\left(\frac{1}{2},\frac{1}{2}\right)$.

Let $l_2 = \frac{1}{2}$ and consider a deviation at some $l_1^{d^{min}} \in \left[0,\frac{1}{2}\right)$. Deviation incentives of firm 2 follow in an analogous manner by relabelling firms. Deviation profits of firm

1 are given by $\Pi_{1\,NB}^{d^{min}} = \left(f_{1\,NB}^{*} + a - c\right)\left(\frac{l_1^{d^{min}}}{2} + \frac{1}{4} - \frac{f_{1\,NB}^{*} - f_{2\,NB}^{*}}{2t\left(\frac{1}{2} - l_1^{d^{min}}\right)}\right)$ if $l_1^{d^{min}}$ is such that the price equilibrium is of type *NB*; that is if $l_1^{d^{min}} \in \left[0, -1 + \frac{\sqrt{9 - \frac{12\left(\underline{f} + \bar{\bar{a}}(1-q) - c\right)}{t}}}{2}\right) = \mathrm{l}_{1\,NB}^{d^{min}}$. Differentiating with respect to $l_1^{d^{min}}$ yields that $\frac{d\Pi_{1\,NB}^{d^{min}}}{dl_1^{d^{min}}} = \left(f_{1\,NB}^{*} + \bar{\bar{a}}(1-q) - c\right)\frac{-\left(l_1^{d^{min}} + \frac{1}{2}\right)}{2\left(\frac{1}{2} - l_1^{d^{min}}\right)} < 0$. This implies that deviation profits have a local maximum at $l_1^{d^{min}} = 0$, thus maximum deviation profits are given by $\Pi_{1\,NB}^{d^{min}} = \frac{25t}{144}$. Comparing with Π_{min} gives that a deviation at $l_1^{d^{min}} = 0$ is profitable if $\frac{25t}{144} > \frac{\underline{f} + \bar{\bar{a}}(1-q) - c}{2} \Longrightarrow \underline{f} + \bar{\bar{a}}(1-q) - c < \frac{25t}{72}$.

Deviation profits are given by $\Pi_{1\,BB}^{d^{min}} = \left(\underline{f} + \bar{\bar{a}}(1-q) - c\right)\left(\frac{l_1^{d^{min}}}{2} + \frac{1}{4}\right)$ if $l_1^{d^{min}}$ is such that the price equilibrium is of type *BB*; that is if $l_1^{d^{min}} \in \left(1 - \frac{\sqrt{\frac{4\left(\underline{f} + \bar{\bar{a}}(1-q) - c\right)}{t} + 1}}{2}, \frac{1}{2}\right) = \mathrm{l}_{1\,BB}^{d^{min}}$. Differentiating with respect to $l_1^{d^{min}}$ yields that $\frac{d\Pi_{1\,BB}^{d^{min}}}{dl_1^{d^{min}}} = \frac{\left(\underline{f} + \bar{\bar{a}}(1-q) - c\right)}{2} > 0$. Thus, deviating to any $l_1^{d} \in \mathrm{l}_{1\,BB}^{d^{min}}$ is less profitable than choosing $l_1 = 0$.

Deviation profits are given by $\Pi_{1\,AB1}^{d^{min}} = \left(\underline{f} + \bar{\bar{a}}(1-q) - c\right)\left(\frac{l_1^{d^{min}}}{2} + \frac{1}{4} - \frac{\underline{f} - f_{2\,AB1}^{*}}{2t\left(\frac{1}{2} - l_1^{d^{min}}\right)}\right)$ if $l_1^{d^{min}}$ is such that the price equilibrium is of type *AB1*; that is if $l_1^{d^{min}} \in$

$\left[-1+\frac{\sqrt{9-\frac{12\left(\underline{f}+\bar{\bar{a}}(1-q)-c\right)}{t}}}{2}, 1-\frac{\sqrt{\frac{4\left(\underline{f}+\bar{\bar{a}}(1-q)-c\right)}{t}+1}}{2}\right) = \mathrm{l}_{1\,AB1}^{d^{min}}$. Differentiating with respect to $l_1^{d^{min}}$ yields that $\frac{d\Pi_{1\,AB1}^{d^{min}}}{dl_1^{d^{min}}} = \left(\underline{f}+\bar{\bar{a}}(1-q)-c\right)\left(\frac{1}{4}-\frac{\left(\underline{f}+\bar{\bar{a}}(1-q)-c\right)}{4t\left(\frac{1}{2}-l_1^{d^{min}}\right)^2}\right)$. First and second order conditions yield a local maximum at $l_1^d = \frac{1}{2}-\sqrt{\frac{\underline{f}+\bar{\bar{a}}(1-q)-c}{t}}$, while lies inside $\mathrm{l}_{1\,AB1}^{d^{min}}$ if $\underline{f}+\bar{\bar{a}}(1-q)-c > \frac{9t}{16}$ and lies to the left of $\mathrm{l}_{1\,AB1}^{d^{min}}$ otherwise. If $\frac{1}{2}-\sqrt{\frac{\underline{f}+\bar{\bar{a}}(1-q)-c}{t}}$ lies inside $\mathrm{l}_{1\,AB1}^{d^{min}}$ then the maximum deviation profits at $l_1^d \in \mathrm{l}_{1\,AB1}^{d^{min}}$ are given by $\frac{\left(\underline{f}+\bar{\bar{a}}(1-q)-c\right)\left(\frac{3\sqrt{\left(\underline{f}+\bar{\bar{a}}(1-q)-c\right)t}}{2}-\left(\underline{f}+\bar{\bar{a}}(1-q)-c\right)\right)}{2\sqrt{\left(\underline{f}+\bar{\bar{a}}(1-q)-c\right)t}}$, which is greater than Π_{min} if $\underline{f}+\bar{\bar{a}}(1-q)-c<\frac{t}{4}$. But $\frac{t}{4}<\frac{9t}{16}$. Thus, if $\frac{1}{2}-\sqrt{\frac{\underline{f}+\bar{\bar{a}}(1-q)-c}{t}}$ lies inside $\mathrm{l}_{1\,AB1}^{d^{min}}$, any deviation at $l_1^{d^{min}} \in \mathrm{l}_{1\,AB1}^{d^{min}}$ is unprofitable. If $\frac{1}{2}-\sqrt{\frac{\underline{f}+\bar{\bar{a}}(1-q)-c}{t}}$ lies to the left of $\mathrm{l}_{1\,AB1}^{d^{min}}$ then the firm 1 is better off deviating at some $l_1^{d^{min}} \in \mathrm{l}_{1\,NB}^{d^{min}}$, as $\frac{d\Pi_{1\,AB1}^{d^{min}}}{dl_1^{d^{min}}} < 0$ for all $l_1^d \in \mathrm{l}_{1\,AB1}^{d^{min}}$.

Finally, there does not exist an $l_1^{d^{min}}$ such that the price equilibrium is of type *AB2*. Therefore, there exists a minimum-differentiation equilibrium if $\underline{f}+\bar{\bar{a}}(1-q)-c \geq \frac{25t}{72}$. ■

Proof of Proposition 3.

(i) By Proposition 2, a maximum-differentiation equilibrium exists if $\frac{t}{2} \geq \underline{f}+\bar{\bar{a}}(1-q)-c$. Since lower values of $\bar{\bar{a}}$ and higher values of q increase the parameter space for which this inequality holds, a market policy that lowers $\bar{\bar{a}}$ or increases q cannot shift the market away from the maximum-differentiation equilibrium.

(ii-a) At a minimum-differentiation equilibrium with $q = q_{old}$, the average consumer pays an total average price $p^{min} = \underline{f} + \bar{\bar{a}}(1 - q_{old})$. At a maximum-differentiation equilibrium with $q = q_{new}$, the average consumer pays a total average price $p^{max} = c + t$. Thus, the average consumer is better off after a disclosure policy if $\underline{f} + \bar{\bar{a}}(1 - q_{old}) > c + t$.

(ii-b) At a minimum-differentiation equilibrium with $q = q_{old}$, the average sophisticated consumer pays an total price $p_s^{min} = \underline{f}$. At a maximum-differentiation equilibrium with $q = q_{new}$, the average sophisticated consumers pays a total price $p_s^{max} = c + t - \bar{\bar{a}}(1 - q_{new})$. Thus, the average sophisticated consumer is better off after a disclosure policy if $\underline{f} > c + t - \bar{\bar{a}}(1 - q_{new})$. But by part *(b)* of Proposition 2, a maximum-differentiation cannot exist for a q_{new} that satisfies this inequality. This the average sophisticated consumers is always made worse off by such a market policy.

(ii-c) At a minimum-differentiation equilibrium with $q = q_{old}$, the average naïve consumer pays an total price $p_n^{min} = \underline{f} + \bar{\bar{a}}$. At a maximum-differentiation equilibrium with $q = q_{new}$, the average naïve consumer pays a total price $p_n^{max} = c + t + \bar{\bar{a}}q_{new}$. Thus, the average naïve consumer that remains uninformed after a disclosure policy is better off after the policy if $\underline{f} + \bar{\bar{a}} > c + t + \bar{\bar{a}}q_{new}$. But, again, by part *(b)* of Proposition 2, a maximum-differentiation cannot exist for q_{new} that satisfies this inequality. Thus, the average uninformed naïve consumer is always made worse off by such a market policy.

(ii-d) At a minimum-differentiation equilibrium with $q = q_{old}$, the average naïve consumer pays an total price $p_n^{min} = \underline{f} + \bar{\bar{a}}$. At a maximum-differentiation equilibrium with $q = q_{new}$, the average sophisticated consumer pays a total price $p_s^{max} = c + t - \bar{\bar{a}}(1 - q_{new})$. Thus, the average naïve consumer that becomes informed after a disclosure policy is better off after the policy if $\underline{f} + \bar{\bar{a}} > c + t - \bar{\bar{a}}(1 - q_{new})$.

(iii-a) At a minimum-differentiation equilibrium with $\bar{\bar{a}} = \bar{\bar{a}}_{old}$, the average consumer pays an total average price $p^{min} = \underline{f} + \bar{\bar{a}}_{old}(1 - q)$. At a maximum-differentiation equilibrium with $\bar{\bar{a}} = \bar{\bar{a}}_{new}$, the average consumer pays a total average

price $p^{max} = c + t$. Thus, the average consumer is better off after a price control policy if $\underline{f} + \bar{\bar{a}}_{old}(1-q) > c + t$.

(iii-b) At a minimum-differentiation equilibrium with $\bar{\bar{a}} = \bar{\bar{a}}_{old}$, the average sophisticated consumer pays a total price $p_s^{min} = \underline{f}$. At a maximum-differentiation equilibrium with $q\bar{\bar{a}} = \bar{\bar{a}}_{new}$, the average sophisticated consumers pays a total price $p_s^{max} = c + t - \bar{\bar{a}}_{new}(1-q)$. Thus, the average sophisticated consumer is better off after a disclosure policy if $\underline{f} > c + t - \bar{\bar{a}}_{new}(1-q)$. But by part *(b)* of Proposition 2, a maximum-differentiation cannot exist for an $\bar{\bar{a}}_{new}$ that satisfies this inequality. This the average sophisticated consumers is always made worse off by such a market policy.

(iii-c) At a minimum-differentiation equilibrium with $\bar{\bar{a}} = \bar{\bar{a}}_{old}$, the average naïve consumer pays a total price $p_n^{min} = \underline{f} + \bar{\bar{a}}_{old}$. At a maximum-differentiation equilibrium with $\bar{\bar{a}} = \bar{\bar{a}}_{new}$, the average naïve consumer pays a total price $p_n^{max} = c + t + \bar{\bar{a}}_{new}q$. Thus, the average naïve consumer is better off after the policy if $\underline{f} + \bar{\bar{a}}_{old} > c + t + \bar{\bar{a}}_{new}q$. ■

Chapter 4

Paying for Desirable Add-ons in a Market with Endogenous Product Differentiation

4.1 Introduction

In many markets, consumers may pay extra to use add-on services that complement their primary purchase. For example, personal current account holders may pay to use overdraft facilities; mobile phone users that exhaust their allowance may pay overage charges to make more phone calls; air travellers may buy on-board meals; car owners may pay for maintenance. Due to behavioural biases such as overconfidence and lack of self-control, consumers often miscalculate their future demand for such services while shopping around for the primary good. As a result, they purchase add-ons more often[42] than they anticipate.

While firms may want to exploit consumers' short-sightedness by setting high fees for add-on services, exploitation is undermined by competition forces. In the absence

[42] See, for example, Office of Fair Trading (2008), par. 4.68, 4.69 on personal current accounts and Federal Communications Commission (2010) on phone plans.

of constraints, firms compete more fiercely for profitable biased consumers, lowering advertised prices possibly below marginal cost and dissipating profits from the unanticipated fees[43]. However, competitive pressure can, in some cases, be inadequate. In particular, firms can retain some excess profit if a lower bound[44] on upfront prices prevents full dissipation. Past literature has studied how retaining some excess profits due to a price floor affects firm behaviour in terms of pricing strategy (Ellison, 2005), aftermarket monopolization efforts (Miao, 2010), shrouding (Heidhues *et al*., 2017), and R&D expenditure (Heidhues *et al*., 2016).

In this chapter we study how the retention of profits from the sale of desirable[45] add-on services interacts with the degree of horizontal product differentiation[46] in the primary market.

In particular, we have two aims. The first is to investigate the effect of exogenously given product differentiation on equilibrium price structure and on the ability of firms to retain extra profits from unanticipated add-on fees. We show that a price equilibrium generally exhibits low upfront prices and high add-on fees, but is not necessarily in pure strategies. The degree of product differentiation critically influences how much extra profits firms can retain.

Our second aim is to explore the levels of product differentiation that endogenously arise in such markets. We show that firms face a trade-off between softening competition in the primary market by offering horizontally differentiated products, and retaining larger profits from unanticipated fees by offering products that are closer substitutes. This trade-off can establish an equilibrium in which profits from

[43] Gabaix and Laibson (2006) contains such a profit irrelevance result regarding unanticipated desirable add-ons. A similar dissipation argument appears in models where firms have monopoly power over an aftermarket due to switching costs (see, for example, Lal and Matutes, 1994; Shapiro, 1995).

[44] For a careful discussion on why price-cutting may be constrained see section 3.1 in Chapter 3. See also sections 2.2 and 2.3 in Heidhues and Koszegi (2018).

[45] We define as desirable the add-on services that can be sold to rational consumers at prices above cost.

[46] Horizontal differentiation may be real or spurious (Spiegler, 2006; Tremblay and Polasky, 2002). We model it in a traditional way, as a choice of product characteristics over which consumers have heterogeneous preferences.

unanticipated fees are eroded, and two equilibria in which some of these profits are retained. Consumers may be better off in the latter.

Our work focuses on markets in which firms sell desirable add-ons at prices that both sophisticated and myopic consumers accept.[47] This feature may apply to markets with sufficiently large population of sophisticated types. It may also apply to markets in which sophisticates' outside option is very costly.

We analyse a simultaneous-move two-stage duopoly model in which firms first choose locations on a Hotelling line and then choose prices. To more simply model a market in which each firm offers a primary good and a desirable add-on service we assume that each firm sells a unique, indivisible product, and charges two separate price components: an upfront price, paid by consumers upon purchase, and an add-on fee, paid by consumers at a later period unless an avoidance cost is incurred. In real world terms, that avoidance cost represents the cost of an outside option.

Not all consumers understand that the add-on fee is avoidable only at a cost. Sophisticated types do so, but naïve types believe they will never pay the add-on fee regardless of whether they pay for avoidance. In real terms, the naïve types erroneously believe that the add-on service is provided for free.

To capture the idea that competition does not always erode profits completely, we assume a lower bound on upfront prices in the form of a price floor. The degree of horizontal differentiation is crucial to whether firms can enjoy any excess profits, as a binding price floor constraint requires products to be sufficiently close substitutes. This creates two conflicting incentives for the firms' location choice: (i) locate far apart to soften competition and enjoy higher upfront prices, but do not take advantage of the price floor constraint, or (ii) locate closer together to compete more fiercely for customers, while retaining some profits from the add-on fees due to the price floor. The relative strength of these incentives depends non-monotonically on interfirm distance. Our model reveals the locations pairs were firms' relocation incentives converge.

[47] Chapter 3 investigates the interaction between add-on profits and horizontal differentiation in markets where add-ons services are unwanted and purchased only by myopic consumers.

In contrast with traditional Hotelling models (Hotelling, 1929; d'Aspremont *et al.*, 1979; Neven, 1985) in which fixed prices lead to minimum differentiation and variable prices lead to maximum differentiation, we find that the combination of consumer naïvete and a relevant price floor can generate maximum, minimum or intermediate differentiation in equilibrium. Intermediate differentiation is unambiguously preferred in terms of average consumer travel costs. Consumers are worse off in the maximum-differentiation equilibrium despite firms dissipating all extra profits.

Our model advises caution when considering regulatory interventions in markets that feature unexpected add-on fees. The effectiveness of interventions depends on the ex-ante degree of product differentiation in the market and on whether firms respond by choosing different product locations. While modest interventions can be beneficial, excessive interventions that can shift the market to a different equilibrium type can lead to losses in both social and consumer surplus.

We contribute to growing study of markets with behavioural consumers (Heidhues and Koszegi, 2018; Grubb, 2015a) by investigating how the firms' choice of product differentiation is affected by their desire to exploit consumer mistakes. We show that, compared to a rational benchmark, exploitative firm behaviour can increase total surplus by improving product matching. Other works show that exploitative firm behaviour can create inefficiencies. Such inefficiencies include excessive consumption of the primary good (Heidhues and Koszegi, 2010; 2015), waste of effort and resources, or altering of consumption patterns in order to avoid future payments (Gabaix and Laibson, 2006; Grubb, 2015b), prevalence of low quality products (Michel, 2018; Gamp and Krahmer, 2018; Heidhues *et al.*, 2017) and redirection of resources away from value-enhancing R&D (Heidhues *et al.*, 2016).

By modelling a market with desirable add-ons whose demand is miscalculated by some consumers, our work contributes to the discussion on whether sophisticated "shoppers" can improve the deals available to all consumer types (Varian, 1980; Armstrong and Chen, 2009; Armstrong, 2015). We show that the opposite is possible; the presence of naïve types can establish equilibria that benefit both types but are unavailable in markets populated solely by sophisticates.

On studying the interplay between market transparency and product differentiation incentives, our work is most closely related to Schultz (2004). For a certain parameter space our model predicts similar firm behaviour to that in Schultz's work. However, our inclusion of both salient prices and unanticipated fees establishes more than a single market outcome. We find that increasing transparency may incentivise firms to offer more, rather than less, differentiated products.

Finally, our policy implications place our paper alongside works that investigate the benefits and unintended consequences of regulatory interventions in markets with naïve consumers (e.g. Grubb, 2015b; Armstrong *et al.* (2009); Kosfeld and Schuwer, 2017). We contribute to this literature by revealing that interventions may influence the degree of product differentiation in the market, leading to losses in both consumer and total surplus.

The rest of the paper is organised as follows. Section 4.2 sets up our model. Section 4.3 presents our equilibrium analysis. In section 4.3.1 we characterise the price equilibrium for fixed symmetric locations. In section 4.3.2 we study firms' relocation incentives and present the location pairs towards which these incentives converge. In Section 4.4 we study welfare and policy implications. Section 4.5 concludes.

4.2 Model

Consider two firms located at l_1, l_2 on a Hotelling line between 0 and 1, and without loss of generality, let $l_1 \leq l_2$. Each firm produces a homogenous good of utility v at cost c, $c < v$. Assume that v is large enough that the market is always fully covered.

Consumers are uniformly distributed on the interval $[0,1]$ and have mass 1. Each consumer has unit demand and may purchase from at most one firm. For a consumer located at $\chi \in [0,1]$, purchasing from firm i entails a distance cost $t(\chi - l_i)^2$.

Purchasing product i entails two price components: an upfront price $f_i \in [c, +\infty)$, where c acts as an exogenous price floor[48], and an add-on fee $a_i \in (-\infty, +\infty)$. Both price components of each firm are known to all consumers before purchase.

[48] The idea that the marginal cost constitutes a price floor is natural in markets with marginal costs very close to zero. Intuitively, a price below zero is unprofitable if it attracts a large number of extremely

In period 1 each consumer purchases from her firm of choice by paying f_i. In period 2, each consumer that has purchased in period 1 from firm i learns she will have to pay the add-on fee, a_i, in period 3. Every consumer pays a_i in period 3, unless that consumer incurs an early-avoidance cost $\bar{a}$ in period 1, or a late-avoidance cost $\bar{\bar{a}}$ in period 2.[49] Assume $\bar{\bar{a}} > \bar{a} > 0$.[50]

Consumers come in two types, depending on their belief about whether they will have to pay the add-on fee, a_i. Proportion q of consumers are sophisticated in that they understand that purchase in period 1 comes with a charge a_i in period 3. Therefore, they will pay the early-avoidance cost $\bar{a}$ in period 1 if they observe $a_i > \bar{a}$. Proportion $1 - q$ of consumers are naïve in that they make a purchase in period 1 believing they will not be charged a_i in period 3. In period 2 they learn they will be charged a_i. Therefore, they pay the late-avoidance cost $\bar{\bar{a}}$ in period 2 if $a_i > \bar{\bar{a}}$. Both types are uniformly distributed on the line. As a tie-breaker, assume that consumers always choose to pay the fee if they are indifferent between paying and avoiding. To ensure that firms prefer to collect add-on fees from both consumer types we assume that $\bar{a} >$

sophisticated consumers that earn the subsidy while costlessly avoiding any add-on fees. If arbitrage is possible then setting upfront price below c may be a dominated strategy even for markets with relatively high marginal costs. This is because it may allow a rival to purchase below cost and resale, effectively gaining a cost advantage. Finally, firms may have preference against below-cost prices to avoid accusations of predatory pricing. See sections 2.2, 2.3 in Heidhues and Koszegi (2018) for a range of works that study markets with price floors.

[49] For example, a consumer purchasing a mobile contract can avoid overage charges in two ways. The first way is by paying attention to her balance and making only the most valuable calls. $\bar{a}$ represents the sum of the necessary attention cost and the utility loss from forfeiting marginal calls. The second way is by stopping using her phones once she is alerted that overage charges apply. $\bar{\bar{a}}$ represents the utility loss from forfeiting all calls after the alert.

[50] The assumption of a positive early avoidance cost ($\bar{a}$) equates to the notion that add-ons are desirable, as it implies that consumers prefer not to avoid add-on consumption if the price of the add-on is not above $\bar{a}$.

$\bar{\bar{a}}(1-q)$.[51] [52] We also assume $\bar{\bar{a}}(1-q) \geq \bar{a}q$ to ensure that a price equilibrium in symmetric locations is always tractable.[53]

The timing is as follows. In stage 1 firms simultaneously choose locations on the line. In stage 2 firm simultaneously set prices. Consumer behaviour consists of three periods. In period 1 consumers observe firm locations and prices f_i, a_i, form expectations about their probability of being charged a_i in period 3, and choose from which firm to purchase and whether or not to incur the early-avoidance cost $\bar{a}$. Sophisticated consumers correctly believe they will be charged a_i in period 3 with probability 1 and choose to avoid the payment by incurring the early-avoidance cost $\bar{a}$ if they observe $a_i > \bar{a}$. Naïve consumers erroneously believe they will not be charged a_i in period 3. In period 2, each naïve consumer that has purchased in period 1 learns that she will be charged a_i in period 3. She chooses to avoid the payment by incurring the late-avoidance cost $\bar{\bar{a}}$ if $a_i > \bar{\bar{a}}$. In period 3, consumers that have not engaged in costly avoidance in the previous periods pay a_i. We solve the game using backward induction looking for symmetric equilibria.

4.3 Equilibrium analysis

4.3.1 Equilibrium price structure for fixed symmetric locations

We first analyse the Nash equilibrium in prices for symmetric exogenously given locations pair (l_1, l_2). An equilibrium may be in pure or in mixed strategies. Our

---

[51] This assumption is satisfied for a sufficiently high proportion of sophisticated consumers. It can be rearranged to $q > \frac{\bar{\bar{a}}-\bar{a}}{\bar{a}}$. Inasmuch as the degree of sophistication is endogenously chosen by consumers, a larger $\bar{\bar{a}} - \bar{a}$ would endogenously bring about a higher q, maintaining the validity of this assumption. A study of consumers' incentives to become sophisticated lies outside the scope of this paper. See Armstrong *et al*. (2009) and Heidhues *et al.* (2018) studies of such incentives.

[52] Chapter 3 models a market in which the add-on component is worthless and easy to avoid – that is, a case in which $\bar{a} = 0$. Specifications with $\bar{a} < \bar{\bar{a}}(1-q)$ are qualitatively similar.

[53] Since $\bar{\bar{a}} > \bar{a}$ this assumption is satisfied if, for instance, $q < \frac{1}{2}$.

analysis proceeds as follows. In Lemmas 1 and 2 we derive the prices that arise in a pure strategy Nash equilibrium, given that such an equilibrium exists. In Lemma 3 we provide conditions under which the derived equilibrium prices result in a pure strategy Nash equilibrium, and conditions under which only an equilibrium in mixed strategies exists. Proposition 1 collects our results.

With full coverage, each firms' demand is determined by the position of the consumer who is indifferent between buying from firm i versus buying from firm j. A sophisticated consumer who makes a purchase in period 1 from firm i realises an associated cost $a_i^s \in \{a_i, \bar{a}, \bar{\bar{a}}\}$, depending on her behaviour. a_i represents her cost if she pays a_i in period 3 to firm i, $\bar{a}$ represents her early-avoidance cost if she chooses to avoid in period 1, and $\bar{\bar{a}}$ represents her late-avoidance cost if she chooses to avoid in period 2. Hence, the location of the indifferent sophisticated consumer is given by x^s, which satisfies $f_1 + a_1^s + t(x^s - l_1)^2 = f_2 + a_2^s + t(x^s - l_2)^2$. Solving for x^s gives $x^s = \frac{l_1+l_2}{2} - \frac{f_1+a_1^s-f_2-a_2^s}{2t(l_2-l_1)}$. In contrast, a naïve consumer makes a purchase expecting an associated cost $a_i^n = 0$. As such, the location of the indifferent naïve consumer is given by x^n, which satisfies $f_1 + t(x^n - l_1)^2 = f_2 + t(x^n - l_2)^2$. Solving for x^n gives $x^n = \frac{l_1+l_2}{2} - \frac{f_1-f_2}{2t(l_2-l_1)}$.

With $l_1 < l_2$, firm 1 (2) serves every consumer to the left (right) of the indifferent consumer. At any symmetric pure strategy equilibrium each firm enjoys demand from both consumer types. Hence, given demand weights q for sophisticated and $1-q$ for naïve consumers, for prices around an equilibrium price pair (f^*, a^*) the demand of firm 1 is given by $D_1 = qx^s + (1-q)x^n$. Accordingly, the demand of firm 2 is given by $D_2 = q(1-x^s) + (1-q)(1-x^n)$.

4.3.1.1 Pure strategy Nash equilibrium

By Lemma 1, equilibrium add-on fees are such that no consumer type prefers to avoid them.

Lemma 1. *If it exists, a pure strategy price equilibrium is characterised by* $a_i^* = a^* \leq \bar{a}$.

Proof: *In appendix C.*

Intuitively, the best a_i for a firm that aims to earn fee revenues from only the naïve types is $\bar{\bar{a}}$, since this is the highest fee for which sophisticated types avoid while naïve types pay. For a firm aiming to earn fee revenues from both consumer types the best fee is some $a_i^* \leq \bar{a}$, as sophisticates avoid any $a_i > \bar{a}$. Assumption $\bar{a} > \bar{\bar{a}}(1-q)$ ensures that, at a pure strategy price equilibrium, firms prefer to earn fees for both types, thus any pricing strategy with $\bar{\bar{a}}$ is strictly dominated.

For prices around an equilibrium price combination (f^*, a^*) demand from either consumer type is given by the second line of equations (1), (2).[54] Making use of symmetry in locations $(l_1 = 1 - l_2)$, we can express each firm's maximisation problem by $\max\limits_{f_i, a_i} \Pi_i \;\; s.t. \;\; f_i \geq c, a_i \leq \bar{a}$, where[55]

$$\Pi_i = (f_i + a_i - c)\left(\frac{1}{2} - \frac{q(a_i - a^*)}{2t(2l_2 - 1)} - \frac{f_i - f^*}{2t(2l_2 - 1)}\right) \qquad (1)$$

The solution to the constrained maximisation problem is presented in Lemma 2. Subscripts in f^*, a^* denote the range of l_2 for which these variables characterise a pure strategy price equilibrium, if such an equilibrium exists. For example, if there exists a pure strategy price equilibrium for $l_2 \in l_2^H$, it is characterised by f_H^*, a_H^*, and Π_H^*.

---

[54] Since, by Lemma 1, $a_i^* \leq \bar{a}$, sophisticated types prefer to pay the add-on fees. As such, we have $a_i^s = a_i$ in the expression for x^s.

[55] Note that the profit function is discontinuous at $a_i = \bar{a}$, since sophisticates avoid any $a_i > \bar{a}$. Then if $a^* = \bar{a}$ equation (1) describes firm i's profits only for a_i slightly below $\bar{a}$. This is not a problem since, by Lemma 1, firm i never plays $a_i > \bar{a}$.

Lemma 2. *Let $l_1 = 1 - l_2$. If it exists, a pure strategy price equilibrium is characterised by:*

(a) $(f_H^*, a_H^*) = (c - \bar{a} + t(2l_2 - 1), \bar{a})$ *and* $\Pi_H^* = \frac{t(2l_2-1)}{2}$ *for* $l_2 \in l_2^H = \left[\frac{1}{2} + \frac{\bar{a}}{2t}, 1\right]$

(b) $(f_I^*, a_I^*) = (c, \bar{a})$ *and* $\Pi_I^* = \frac{\bar{a}}{2}$ *for* $l_2 \in l_2^I = \left(\frac{1}{2} + \frac{q\bar{a}}{2t}, \frac{1}{2} + \frac{\bar{a}}{2t}\right)$

(c) $(f_L^*, a_L^*) = \left(c, \frac{t}{q}(2l_2 - 1)\right)$ *and* $\Pi_L^* = \frac{t(2l_2-1)}{2q}$ *for* $l_2 \in l_2^L = \left[\frac{1}{2}, \frac{1}{2} + \frac{q\bar{a}}{2t}\right]$

(d) *Every* (f^*, a^*) *such that* $f^* + a^* = c + t(2l_2 - 1)$ *if* $q = 1$

Proof: *In appendix C.*

To get an intuition behind why there does not exist a pure strategy price equilibrium with both $f^* > c$ and $a^* < \bar{a}$ recall that the demand is more elastic with respect to upfront prices than with respect to add-on fees. Then, a firm choosing (f_i, a_i) such that $f_i > c$, $a_i < \bar{a}$ can do better by increasing its fee and reducing its upfront price by an equal amount. In that way, the firm enjoys the same mark-up while attracting more naïve consumers (who pay no attention to add-on fees when purchasing). Incentives to back-load are exhausted either when $f_i = c$, as the price floor prevents lower upfront prices, or when $a_i = \bar{a}$, as $a_i > \bar{a}$ results in sophisticates avoiding any fee payments, which is unprofitable by Lemma 1.

The difference in demand elasticity between fees and upfront prices is also the reason why the shape of f^*, a^*changes at different interfirm distance intervals. At any location pair, it is in the firms' best interest to compete primarily on upfront prices while charging the highest possible fees, as a cut on upfront prices attracts both naïve and sophisticated consumers. It is only for location pairs for which the price floor constraint binds – that is, competition on upfront prices is exhausted – that firms have any incentive to compete on fees. These location pairs are described in parts *(b)* and *(c)* of Lemma 2.

To see why the price floor is binding when products are located relatively close together, note that the equilibrium upfront price f^* falls with lower interfirm distance.

This is because competition is more intense the lower the product differentiation in the market. Consequently, the price floor constraint is stricter the closer firms are located.

For $l_2 \in l_2^I$ firms' incentive to compete on upfront prices is exhausted. However, neither do they have an incentive to compete on fees. The reason is that any cut in fees reduces the mark-up from every customer but only attracts a small number of sophisticates. Cutting fees is profitable only if products are located sufficiently close together – for location pairs with $l_2 \in l_2^L$. Due to the quadratic transport costs, the effect of a reduction in a_i increases firm i's demand more the closer products are located. To see this analytically, notice that, for symmetric locations with $f_i = c$, we have $\frac{\partial \Pi_i}{\partial a_i} = \left(\frac{1}{2} - \frac{q(a_i - a^*)}{2t(2l_2-1)}\right) + (a_i)\left(-\frac{q}{2t(2l_2-1)}\right)$. Then, the demand effect of a shift in a_i is stronger than the mark-up effect only if $l_2 < \frac{1}{2} + \frac{q(-\bar{a}+2a_i)}{2t}$. Substituting $a_i = \bar{a}$ gives that cutting fees below $\bar{a}$ can only be profitable if $l_2 < \frac{1}{2} + \frac{q\bar{a}}{2t}$, which is exactly described by interval l_2^L.

Part *(d)* describes the pricing stage in a market with no naïve types. Demand elasticity with respect to upfront prices is equal to demand elasticity with respect to fees since sophisticates correctly anticipate they will pay for both price components. Then the sum of the price components in a symmetric pure strategy equilibrium is identical to the equilibrium price in a standard Hotelling model with a single, salient price. Consequently, the location stage will also be played in an identical way. For the rest of the paper we will ignore this case and analyse the model assuming there is a population of both consumer types in the market.

4.3.1.2 Mixed strategy Nash equilibrium

For a pure strategy price equilibrium to exist it must hold that firms have incentive to both serve and extract a^* from some sophisticated consumers. However, if l_2 is too small this is not the case. Recall that a_L^* falls with lower interfirm distance, while f_L^* is constant at c. Then, since demand from naïve types is independent of a_i, there exist $l_2 \in l_2^L$ for which firms prefer to raise their fees to $\bar{\bar{a}}$ and earn fee revenues only from naïve consumers, over earning fee revenue a_L^* from every customer type. Note that a deviation to $\bar{\bar{a}}$ can only be profitable for $l_2 \in l_2^L$. This is because $a^* = \bar{a}$ for higher interfirm distance, and firms prefers to earn $\bar{a}$ from both types over earning $\bar{\bar{a}}$ from only the naïve types. This critical l_2 below which a deviation to $\bar{\bar{a}}$ is profitable is presented in Lemma 3.

Lemma 3. *Let $l_1 = 1 - l_2$. A pure strategy price equilibrium exists only if $l_2 \geq \frac{1}{2} + \frac{\bar{\bar{a}}q(1-q)}{2t}$. If $l_2 < \frac{1}{2} + \frac{\bar{\bar{a}}q(1-q)}{2t}$ there exists a mixed strategy price equilibrium characterised by $f^* = c$ and a^* randomly drawn from a distribution, with expected profits $\Pi_i = \frac{\bar{\bar{a}}(1-q)}{2}$.*

Proof: *In appendix C.*

The intuition behind the non-existence of a pure strategy price equilibrium follows the logic of Varian (1980). When products are relatively close substitutes, upfront prices are competed down to the price floor. As naïve types pay no attention to add-on fees, they purchase from the firm that is closest to them.[56] Then, if rival's fees are too low, a firm may find it profitable to set a price combination $(c, \bar{\bar{a}})$ and profit from all its naïve customers while abandoning all sophisticates.[57] Interfirm distance is critical for such behaviour to arise. For $l_2 \geq \frac{1}{2} + \frac{\bar{\bar{a}}q(1-q)}{2t}$ no firm has an incentive to focus exclusively the naïve types, as earning a^* from customers of both types is more profitable than earning $\bar{\bar{a}}$ from only naïve types.

In the associated mixed strategy price equilibrium firms compete fiercely on upfront prices, resulting in $f_i^* = c$. Add-on fees, on the other hand, are randomly chosen from a distribution. $\bar{\bar{a}}$ always belongs in that distribution, as it is the optimal fee for a firm that aims to extract the highest possible surplus from its naïve customers.

---

[56] Naïve consumers behave exactly as in Varian (1980) when products are perfect substitutes, in which case they randomly choose from which firm to purchase.

[57] Proof of Lemma 3 shows that a firm which aims to serve only naïve types maximises profits with $(f_i, a_i) = (c, \bar{\bar{a}})$.

Since profits are identical at any price combination in the distribution, expected profits at the mixed strategy price equilibrium are given by $\Pi_i = \frac{\bar{\bar{a}}(1-q)}{2}$. Intuitively, at a price combination $(c, \bar{\bar{a}})$ firm i earns $\bar{\bar{a}}$ from its naïve customers, since they pay both price components, and 0 from any sophisticated customers, since they avoid the add-on fee.[58] Proposition 1 collects our results with respect to the pricing stage.

Proposition 1. (Price equilibria for symmetric fixed locations). *Let* $l_1 = 1 - l_2$.

(a) *For* $l_2 \in l_2^H = \left[\frac{1}{2} + \frac{\bar{a}}{2t}, 1\right]$ *there exists a pure strategy price equilibrium with* $(f_H^*, a_H^*) = (c - \bar{a} + t(2l_2 - 1), \bar{a})$ *and* $\Pi_H^* = \frac{t(2l_2-1)}{2}$.

(b) *For* $l_2 \in l_2^I = \left(\frac{1}{2} + \frac{q\bar{a}}{2t}, \frac{1}{2} + \frac{\bar{a}}{2t}\right)$ *there exist a pure strategy price equilibrium with* $(f_I^*, a_I^*) = (c, \bar{a})$ *and* $\Pi_I^* = \frac{\bar{a}}{2}$.

(c) *For* $l_2 \in l_2^L = \left[\frac{1}{2} + \frac{\bar{\bar{a}}q(1-q)}{2t}, \frac{1}{2} + \frac{q\bar{a}}{2t}\right]$ *there exists a pure strategy price equilibrium with* $(f_L^*, a_L^*) = \left(c, \frac{t}{q}(2l_2 - 1)\right)$ *and* $\Pi_L^* = \frac{t(2l_2-1)}{2q}$

(d) *For* $l_2 \in l_2^M = \left[\frac{1}{2}, \frac{1}{2} + \frac{\bar{\bar{a}}q(1-q)}{2t}\right)$ *there exists a mixed strategy price equilibrium with* $f_M^* = c$, a_M^* *drawn randomly from a distribution, and* $\Pi_M^* = \frac{\bar{\bar{a}}(1-q)}{2}$.

Proposition 1 sums up our findings on how product differentiation, traditionally represented by $t(l_2 - l_1)$ in Hotelling models, is crucial to determining the price structure. To better comprehend its effect, it is useful to decompose product differentiation into two distinct elements: inherent differentiability, denoted by t, and the degree to which firms exploit the inherent differentiability, denoted by interfirm distance $l_2 - l_1$. For ease in exposition we make use of the symmetry in firm locations

[58] At the mixed strategy price equilibrium firms share the naïve types since $f_1^* = f_2^* = c$. Demand from sophisticated types depends on the distribution of add-on fees. Schultz (2005) characterises firm demand in a similar problem.

and denote interfirm distance by l_2, with high (low) l_2 representing large (small) interfirm distance.

Keeping everything else constant, Proposition 1 shows that high interfirm distance ($l_2 \in l_2^H$) establishes a pure strategy price equilibrium in which firms set fees equal to the early-avoidance cost and compete on the upfront price. Intermediate interfirm distance ($l_2 \in l_2^I$) establishes a pure strategy price equilibrium in which fees are equal to the early-avoidance cost and upfront prices are competed down to the price floor. For small interfirm distance ($l_2 \in l_2^L$) firms, unable to compete further on upfront prices, compete for the business of sophisticated types by reducing fees. Finally, for extremely small interfirm distance ($l_2 \in l_2^M$) firms exhaust competition on upfront prices and randomise over fees.

The extent to which changes in interfirm distance correspond to different price structures depends crucially on the relative size of the early avoidance cost $\bar{a}$, the late-avoidance cost $\bar{\bar{a}}$, and the inherent differentiability t.

Everything else constant, higher $\bar{a}$ affects equilibrium price structure in two ways. First, given a binding price floor constraint ($f_i \geq c$), higher $\bar{a}$ reduces the parameter space for which the early-avoidance constraint ($a_i \leq \bar{a}$) is binding. Second, given a binding early-avoidance constraint, higher $\bar{a}$ increases the parameter space for which the price floor constraint is binding. The first effect is obvious, as $\bar{a}$ represent the highest add-on fee that firms can set for which no consumer prefers costly avoidance. To get the intuition behind the second effect recall that firms have an incentive to use profits from add-on fees to subsidise upfront prices.[59] This implies the price floor constraint is more likely to be binding with higher $\bar{a}$, as that results in lower equilibrium upfront prices.

Higher $\bar{a}$ shrinks the parameter space for which $l_2 \in l_2^H$, as it reduces the interfirm distance for which the price floor constraint is slack. For $\bar{a} > t$ there does not exist a parameter combination for which firms set upfront prices above the price floor. Consequently, l_2^H is empty. The parameter space for which $l_2 \in l_2^I$ is affected by $\bar{a}$ in

[59] Subsidisation of upfront prices is common in problems with second-period prices (Shapiro, 1995; Gabaix and Laibson, 2006). The extent of the subsidisation in the context of our model is discussed further below.

two ways. First, the upper bound of l_2^I increases in $\bar{a}$ as the price floor constraint is more easily binding. Second, the lower bound of l_2^I also increases in $\bar{a}$, as the early-avoidance constraint is less easily binding. The net effect is generally positive since $q \leq 1$; that is, the parameter range for which $l_2 \in l_2^I$ increases with $\bar{a}$. l_2^I shrinks as q approaches 1. Intuitively, firms' incentive to compete on fees becomes stronger as the proportion of sophisticated types in the market increases. As a result, the range of l_2 for which firms exhaust competition on upfront prices but do not compete on fees, l_2^I, becomes smaller. Lastly, note that there cannot exist a symmetric location pair on the line that satisfies $l_2 \in l_2^I$ if $\bar{a}q > t$. Intuitively, firms have some incentive to compete on fees at any interfirm distance when inherent differentiability is small, since products are relatively close competitors regardless of l_2. Finally, the parameter space for which $l_2 \in l_2^L$ becomes larger as $\bar{a}$ increases, since the parameter space for which early-avoidance constraint is binding becomes smaller.

A higher late-avoidance cost ($\bar{\bar{a}}$) increases the parameter space for which the price equilibrium is in mixed strategies – the parameter space for which $l_2 \in l_2^M$. Intuitively, higher $\bar{\bar{a}}$ increases the profitability of maximally exploiting naïve types. Accordingly, the parameter space for which $l_2 \in l_2^L$ shrinks, as firm profits in a pure strategy price equilibrium became smaller compared to profits that firms can earn by extracting $\bar{\bar{a}}$ from just the naïve types.

The effect of inherent differentiability (t) on price structure moves in a direction opposite to the effect of $\bar{a}$. Intuitively, higher t increases the extent to which higher interfirm distance affects upfront prices and fees. As a result, it increases the parameter space for which the price floor is slack and the parameter space for which the early-avoidance cost is binding. Higher t also reduces the parameter space for which $l_2 \in l_2^M$. This is because t generally increases the fee revenues that firms enjoy from sophisticated types for given l_2, making a deviation to $\bar{\bar{a}}$ less profitable in comparison. To see this notice that a_L^* is increasing in t.

A direct implication of Proposition 1 regarding the profitability of markets with naïve consumer types and non-salient fees is that large increases in product differentiation increase industry profits, but small increases in product differentiation do not necessarily do so. The former follows directly from the fact that $\Pi_H^*|_{l_2 \in l_2^H} > \Pi_I^*|_{l_2 \in l_2^I} > \Pi_L^*|_{l_2 \in l_2^L} > \Pi_M^*|_{l_2 \in l_2^M}$. To see the latter, notice that, for a small increase in interfirm distance to increase profits, it must be that at least one constraint is slack. In other words, it must be that $l_2 \in l_2^H$ or $l_2 \in l_2^L$. For markets with $l_2 \in l_2^I$, a small increase in l_2 has no effect on profits, as competition is sufficiently intense to bring equilibrium upfront prices equal to the price floor, but not intense enough to incentivise firms to compete on fees. For markets with $l_2 \in l_2^M$ products are so close substitutes that equilibrium upfront prices are competed down to the price floor, while firms randomise over fees. Since expected profits are equal to $\Pi_M^* = \frac{\bar{\bar{a}}(1-q)}{2}$ and the late-avoidance fee is not a function of interfirm distance, we can conclude that l_2 has no effect on Π_M^*.

A crucial problem in markets with naïve consumers and unexpected fees is understanding the extent to which firms enjoy extra profits from such fees; in other words, the extent to which revenues from such fees subsidise lower upfront prices. To better approach this question, it is useful to derive the profits that firms would enjoy in absence of add-on fees – the sophisticated-consumers benchmark. We define the benchmark as the result of each firm solving the unconstrained maximisation problem $\max_{p_i}(p_i - c)\left(\frac{1}{2} - \frac{p_i - p^*}{2t(2l_2-1)}\right)$, where p_i denotes a price combination $f_i + a_i$ and p^* denotes an equilibrium price combination $f^* + a^*$.[60] Maximisation yields profits at the sophisticated-consumers benchmark $\Pi_{bench}^* = \frac{t(2l_2-1)}{2}$. Proposition 1 states that the extent to which firms enjoy profits beyond the sophisticated-consumers benchmark depends largely on how differentiated the products are.

Firms fully erode fee revenues when product differentiation is sufficiently high ($l_2 \in l_2^H$), as illustrated by the fact that f_H^* is reduced by the term $\bar{a}$.[61] Full subsidisation

---

[60] This is the profit maximisation problem in the standard Hotelling model with variable prices and quadratic distance costs. Firms compete on the total price, as sophisticated consumers give equal weights to both price components. A price floor is never binding, as firms can compete on add-on fees if upfront prices cannot be reduced further.

[61] In this case fee revenues are conceptually equivalent to a reduction in marginal costs by $\bar{a}$. In our model the pass-through rate to such implicit cost reduction is 1. For more general settings that host downward sloping demands and variable marginal costs, pass-through rate depends on demand and supply elasticities. For a general treatment of cost pass-through, see Weyl and Fabinger (2013). For an application of cost pass-through analysis to markets with hidden fees, see Agarwal *et al*. (2014).

of upfront prices is permitted due to the price floor constraint being slack even when subsidy $\bar{a}$ is taken into account. The complete erosion of fee revenues results in equilibrium profits that are identical to those in the sophisticated-consumers benchmark. In contrast, subsidisation is only partial for intermediate product differentiation ($l_2 \in l_2^I$), as the price floor constraint does not permit upfront prices to be subsidised below c. As a result, for given locations, firm profits are higher than those in the sophisticated-consumers benchmark. To see this, notice that $\frac{\bar{a}}{2} > \frac{t(2l_2-1)}{2}$ for all $l_2 \in l_2^I$. Partial subsidisation also characterises a market with low product differentiation ($l_2 \in l_2^L$), as upfront prices are constrained by c. However, fee revenues are below $\bar{a}$ per customer, since firms compete on fees in order to attract sophisticated types. As a result, firm profits for $l_2 \in l_2^L$ are lower than profits for $l_2 \in l_2^I$. However, profits for $l_2 \in l_2^L$ are still higher than those in the sophisticated-consumers benchmark. This is because any cut on fees attracts only sophisticated types, giving firms less incentive to compete on fees. Finally, partial subsidisation characterises the price equilibrium also when products are extremely close substitutes ($l_2 \in l_2^M$), as firms play $f_i = c$ with certainty. The extent of retained fee revenues depends on a_M^*, which is randomly drawn. Average profits are higher than profits in the sophisticated-consumers benchmark, since, regardless of the rival's fees, firms can earn $\bar{\bar{a}}$ from every naïve customer.

Figure 4.1 illustrates how profits of firm 2 change with the degree of product differentiation.[62] Profits of firm 1 follow by symmetry. To include every case in Proposition 1, figure 4.1 assumes $t > \bar{a}$. The solid line represents the profit functions described in Proposition 1. Profits are weakly increasing in l_2; strictly increasing for $l_2 \in l_2^H, l_2^L$ and constant for $l_2 \in l_2^I, l_2^M$. The dotted line represents profits at the sophisticated-consumers benchmark, strictly increasing for all l_2. The vertical distance between the solid and the dotted line area represents the additional surplus that firms enjoy due to retained fee revenues. Demonstrably, it is only in sufficiently competitive

[62] Note the similarities between this figure and figure 3.1 in chapter 3. In chapter 3 firms have no incentive to compete on add-on fees, as sophisticated types avoid them while naïve types ignore them. As such, the associated figure 3.1 depicts a simpler relationship between profits and interfirm distance.

markets ($l_2 \in l_2^M, l_2^L, l_2^I$) that firms are benefitting from the non-salience of add-on fees.[63]

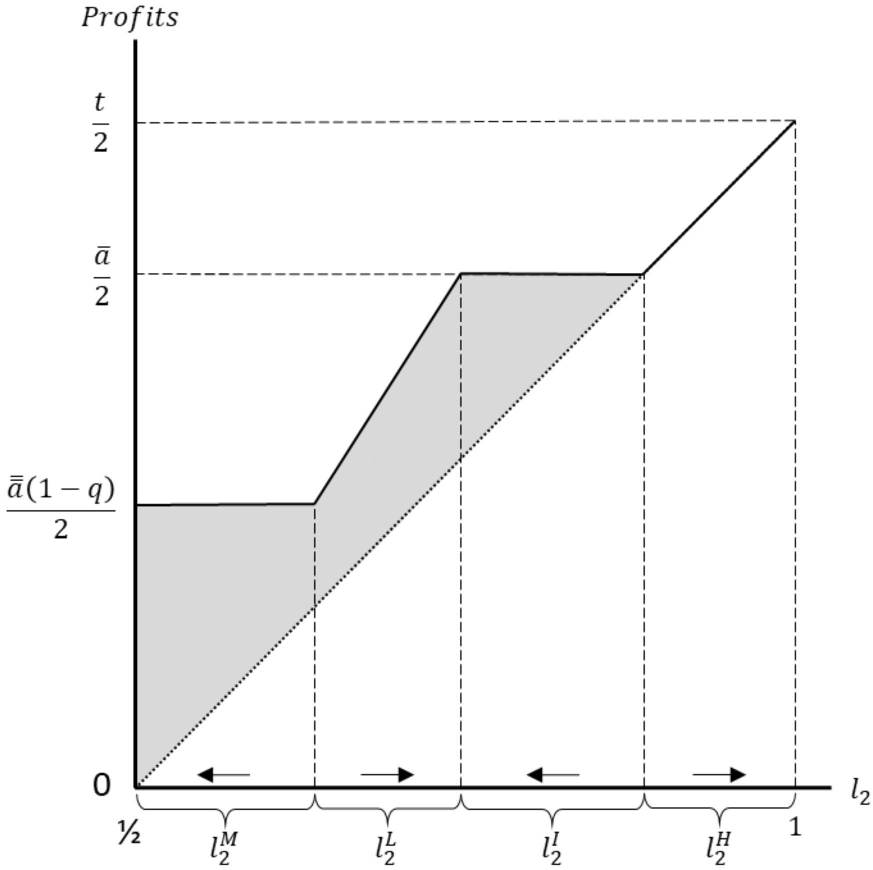

Figure 4.1: Profit and marginal relocation incentives of firm 2 for symmetric locations with $t > \bar{a}$.

Profits and marginal relocation incentives of firm 1 follow by symmetry. The solid line represents profits of firm 2 for symmetric locations with given l_2. The dotted line represents profits of firm 2 at the sophisticated-consumers benchmark. The vertical distance between the solid and the dotted line represents the additional surplus that firm 2 enjoys due to retained fee revenues. Arrows depict marginal relocation incentives of firm 2 for each given symmetric location pair. An inward-pointing arrow means that firm 2 prefers to offer a more homogeneous product; an outward-pointing arrow means that firm 2 prefers to further differentiate its product. Marginal relocation incentives converge towards locations $\frac{1}{2}$, $\frac{1}{2} + \frac{\bar{a}q}{2t}$, and 1, as described in Proposition 2.

[63] Put in the context of primary goods with aftermarket services (Shapiro, 1995; Carlton and Waldman, 2009), this result shows that competition in the primary good needs to be sufficiently intense for firms to profit from aftermarket power.

4.3.2 Local equilibria in locations

For any given equilibrium price strategy, each firm's choice of location on the line depends on two incentives. The first is the incentive to soften competition by offering a differentiated product – the price effect. The second is the incentive to capture a larger market share by offering a product that is a close substitute to the rival's – the demand effect. The relative strength of these effects determines whether firms prefer to locate close to or far apart from their rivals.

In the standard Hotelling model with quadratic distance costs and endogenous prices (d'Aspremont *et al.*, 1979; Neven, 1985) the strength of each incentive changes smoothly as firms locate further apart, since profits functions are continuous and differentiable with respect to locations for any (l_i, l_j). This results in relocation incentives that converge to a unique location pair. In contrast, profit functions in our model are kinked, as the equilibrium price structure changes depending on interfirm distance. As a result, relocations incentives can differ drastically depending on firm locations.

For symmetric location pairs with large interfirm distance ($l_2 \in l_2^H$), a price effect is present, as equilibrium upfront prices are increasing in interfirm distance (l_2). In particular, the price effect dominates the demand effect, so firms prefer to locate further apart. To get an intuition notice that for $l_2 \in l_2^H$ firms behave as in the standard Hotelling model with quadratic distance costs and marginal costs equal to $c - \bar{a}$.

For symmetric location pairs with intermediate interfirm distance ($l_2 \in l_2^I$), a price effect does not arise, as both upfront prices and add-on fees are fixed to c and $\bar{a}$ respectively. As a result, each firm's marginal relocation incentive is dominated by the demand effect. Thus for $l_2 \in l_2^I$ each firm prefers to increase its turf by locating closer to its rival.

For symmetric location pairs with relatively small interfirm distance ($l_2 \in l_2^L$) a price effect is again present, as firms can increase equilibrium add-on fees by locating further apart. This price effect is stronger that the demand effect, resulting in firms preferring to differentiate their products. It is worth noting that the price effect for symmetric location pairs with $l_2 \in l_2^L$ is stronger than that for symmetric location pairs with $l_2 \in l_2^H$; that is, firms profits increase more quickly in interfirm distance for $l_2 \in$

l_2^L than for $l_2 \in l_2^H$. Intuitively, this is because firms with $l_2 \in l_2^L$ compete on fees, which are salient only to sophisticated types, whereas firms with $l_2 \in l_2^H$ compete on upfront prices, which are salient to both consumer types. As a result, firms with $l_2 \in l_2^L$ can take better advantage of larger interfirm distance, since an increase in fees in associated by a smaller loss in demand compared to an increase in upfront prices. It follows that firms have a stronger incentive to increase interfirm distance when $l_2 \in l_2^L$.[64]

Finally, for symmetric location pairs that result in a mixed strategy price equilibrium ($l_2 \in l_2^M$), each firm's expected profits are equal to its profits if it ignores all sophisticated types and extracts $\bar{\bar{a}}$ from its naïve customers. Since $\bar{\bar{a}}$ is independent of interfirm distance, a price effect does not arise. As a result, each firm prefers to locate closer to its rival in order to attract more naïve types.

The arrows in Figure 4.1 illustrate the marginal relocation incentives of firm 2 for each given symmetric location pair. An inward-pointing arrow means that firm 2 prefers to offer a more homogeneous product; an outward-pointing arrow means that firm 2 prefers to further differentiate its product. Marginal relocation incentives of firm 1 follow by symmetry.

Marginal relocation incentives are exhausted – that is, deviations to nearby locations are unprofitable – for the locations pairs described in Proposition 2. These location pairs constitute local Nash equilibria. We restrict our attention to the case in which $t > \bar{a}$, in order to avoid parameter specifications for which some interfirm distance sets $(l_2^H, l_2^I, l_2^L, l_2^M)$ are infeasible. Notation *max*, *min* and *int* indicates the local equilibrium that each variable characterises.

[64] This outcome is consistent with the finding of Schultz (2004) that increasing market transparency (i.e. increasing the proportion of sophisticated types) weakens firms' incentive to locate further apart.

Proposition 2. (Symmetric local equilibrium). *Let* $t \geq \bar{a}$.

(a) *There exists a symmetric maximum-differentiation local equilibrium with* $(l_1^{max}, l_2^{max}) = (0,1)$, $(f_{max}, a_{max}) = (c - \bar{a} + t, \bar{a})$ *and* $\Pi_{max} = \frac{t}{2}$.

(b) *There exists a symmetric intermediate-differentiation local equilibrium with* $\left(l_1^{int}, l_2^{int}\right) = \left(\frac{1}{2} - \frac{q\bar{a}}{2t}, \frac{1}{2} + \frac{q\bar{a}}{2t}\right)$,$(f_{int}, a_{int}) = (c, \bar{a})$, *and* $\Pi_{int} = \frac{\bar{a}}{2}$.

(c) *There exists a symmetric minimum-differentiation local equilibrium with* $\left(l_1^{min}, l_2^{min}\right) = \left(\frac{1}{2}, \frac{1}{2}\right)$, $f_{min} = c$, a_{min} *drawn randomly from* $\left[\bar{\bar{a}}\frac{(1-q)}{1+q}, \bar{a}\right) \cup \{\bar{\bar{a}}\}$, *and* $\Pi_{min} = \frac{\bar{\bar{a}}(1-q)}{2}$.

No other local equilibrium exists.

Proof: *In appendix C.*

Proposition 2 provides important insight regarding the equilibrium structure of markets with naïve consumers and add-on fees that are paid by both consumer types. It unveils that such markets are expected to operate in three equilibrium types.

The first type is described by the maximum-differentiation local equilibrium. To get an intuition behind the existence of this local equilibrium notice that (l_1^{max}, l_2^{max}) describes a symmetric location pair with $l_2 \in l_2^H$. Consequently, firms have an incentive to further differentiate their products, but any further differentiation is impossible as firms already enjoy maximum interfirm distance. Since firm locations satisfy $l_2 \in l_2^H$, this local equilibrium is characterised by upfront prices above the price floor, and add-on fees equal to the consumers' early-avoidance cost. Firm profits are equal to the sophisticated-consumers benchmark $\left(\Pi_{max} = \Pi_{bench}^*|_{l_2=1}\right)$, as firms dissipate all fee revenues by subsidising upfront prices.

The second type is described by the intermediate-differentiation local equilibrium. To get an intuition behind the existence of this local equilibrium notice that l_2^{int} lies between l_2^L and l_2^I. This means that symmetric location pairs with interfirm distance slightly smaller than or equal to $\left(l_1^{int}, l_2^{int}\right)$ are characterised by $l_2 \in l_2^L$, so firms have an incentive to locate further apart in order to soften competition on the add-on fees. Symmetric locations with interfirm distance slightly larger than $\left(l_1^{int}, l_2^{int}\right)$ are characterised by $l_2 \in l_2^I$, so firms have an incentive to locate closer together in order to increase their turf. As a result, marginal relocation incentives from both sides converge towards $\left(l_1^{int}, l_2^{int}\right)$. Location pair $\left(l_1^{int}, l_2^{int}\right)$ represents the minimum interfirm distance for which firms have no incentive to reduce add-on fees below the early-avoidance cost. As such, the intermediate-differentiation local equilibrium is characterised by upfront prices equal to the price floor and add-on fees equal to the consumers' early-avoidance cost. It follows that firms in the intermediate-differentiation local equilibrium retain some fee revenues, as the price floor constraint prevents full dissipation. Thus, firm profits are higher than the sophisticated-consumers benchmark ($\Pi_{int} > \Pi_{bench}^*|_{l_2=\frac{1}{2}+\frac{\bar{a}q}{2t}}$). However, this does not imply that an intermediate differentiation structure results in higher profits than a maximum differentiation structure. On the contrary, the more intense competition on upfront prices due to the limited product differentiation results in lower profits overall.

Interfirm distance in the intermediate-differentiation local equilibrium depends crucially on the relative values of $\bar{a}$, q, t. Everything else constant, a higher early-avoidance cost ($\bar{a}$) increases product differentiation. This is because it permits firms to increase fees further by locating further apart. A high proportion of sophisticated consumers (q) works in the same direction. Since it is only the sophisticated types that respond to differences in fees, increasing the proportion of sophisticated types strengthens each firm's incentive to soften competition on fees by differentiating its product. Higher inherent differentiability (t) stifles firms' incentive to locate marginally further apart. Intuitively, higher t increases the effect that a small increase in interfirm distance has on fees. As a result, the early-avoidance constraint $a_i \leq \bar{a}$ binds with less product differentiation.

The third type is described by the minimum-differentiation local equilibrium. To get an intuition behind its existence notice that $\left(l_1^{min}, l_2^{min}\right)$ describes a symmetric location pair with $l_2 \in l_2^M$, implying that each firm has individual incentive to offer a product that is a closer substitute to its rival's product. Since firms located at $\left(l_1^{min}, l_2^{min}\right)$ already offer products that are perfect substitutes, any relocation to increase substitutability is not feasible. The minimum-differentiation local equilibrium is characterised by upfront prices equal to the relevant price floor, and temporal dispersion in add-on fees. In particular, firms randomise between fee values that all consumers prefer to pay and fee values that sophisticated consumers prefer to avoid. To see why there is a gap in the fee support notice that no firm has incentive to set fees between $\bar{a}$ and $\bar{\bar{a}}$, as that results in sophisticated types preferring costly avoidance. Therefore, firms are better off with fees equal to $\bar{\bar{a}}$ and maximise fee revenues from naïve customers. Firms in the minimum-differentiation local equilibrium retain some fee revenues due to the binding price floor, and earn profits that are, on average, higher than the sophisticated-consumers benchmark ($\Pi_{min} > \Pi^*_{bench}|_{l_2=\frac{1}{2}}$). However, as competition pushes fees below the early-avoidance cost, profits are lower than those in an intermediate-differentiation local equilibrium.

In Proposition 2 we restricted our attention to the case in which $t \geq \bar{a}$. Other parameter specifications do not necessarily permit all three local equilibria. The maximum-differentiation local equilibrium does not exist for $t < \bar{a}$, as l_2^H is empty. Intuitively, firms have weaker incentive to take use product differentiation to soften competition on the upfront prices if the inherent product differentiability (t) too small to permit upfront prices above the price floor. Similarly, the intermediate-differentiation local equilibrium does not exist for inherent product differentiability small enough to satisfy $t < \bar{\bar{a}}q(1-q)$, as this results in $l_2 \in l_2^M$ for any symmetric location pair. Intuitively, such low value of t implies that products are close substitutes regardless of interfirm distance, so firms play a mixed strategy price equilibrium at any symmetric location pair. As such, at any symmetric location pair firms prefer to locate closer together. Finally, the intermediate-differentiation local equilibrium, if it exists, is not characterised by the same equilibrium fees if inherent product differentiability satisfies $t \leq \bar{a}q$. Intuitively, $t \leq \bar{a}q$ implies that products cannot be differentiated enough for equilibrium fees to be reach the early-avoidance cost; in other words the upper bound of l_2^L is greater than 1. Consequently, firms' incentive to differentiate their products for any location pair with $l_2 \in l_2^L$ is exhausted when $l_2 = 1$, resulting in maximum product differentiation and equilibrium fees equal to $\frac{t}{q}$, below $\bar{a}$.

The following section discusses the effectiveness of certain types of regulation in each market type and compare the local equilibria of Proposition 2 in terms of welfare.

4.4 Welfare and policy implications

Given the inability of competition forces to fully shield naïve or unaware consumers from exploitation, regulatory bodies have taken up the role of monitoring markets, often imposing remedies to improve market performance.[65] Some market remedies aim at improving consumers' ability to evaluate products, incentivising firms to offer better deals. For example, in order to better inform prospective mutual funds investors, Securities and Exchange Commission's requires funds' prospectuses to include a summary section with key information.[66] Other remedies aim to directly influence the deals offered to consumers by placing a limit to non-salient add-on fees (or by improving quality). Financial Conduct Authority's price cap on payday loans is an example of such an outcome control remedy, imposing a total cost cap and limiting various charges.[67] In this section we explore the extent to which similar interventions can be used to improve market outcomes in the context of our model.

We investigate the effectiveness of interventions under two welfare standards: average consumer surplus and total surplus. To study changes in average consumer surplus, we explore how interventions affect firm profits and average consumer travel costs. Firm profits fully describe the surplus transfers between consumers and firms; an increase in profits represents and equal and opposite change to the surplus of the average consumer. Naturally, the optimal outcome for consumers in this dimension is a market with zero profits. Average consumer travel costs represent how well, on average, products match consumer tastes. Increasing average travel costs corresponds to a worsening in matching, as it increases the real or psychological cost that the average consumers need to travel for their preferred product. To study changes in total surplus we explore how interventions affect average travel costs alone, as that fully describes the deadweight loss due to the imperfect matching between firm locations and consumer tastes.

---

[65] See Fletcher (2016) for a review on the effectiveness of various types of market remedies.

[66] https://www.sec.gov/rules/final/2009/33-8998.pdf

[67] https://www.fca.org.uk/static/documents/policy-statements/ps14-16.pdf

Our framework is not suited to study other sources of welfare loss that may be relevant to regulators, such as avoidance costs and participation distortions. Their study is left for future work.

4.4.1 Policies that maintain the local equilibrium

We first discuss the desirability of policies that do not introduce a shift in the local equilibrium. We investigate the effectiveness of three categories of policies. The first category is policies that educate consumers about the relevance of add-on fees, resulting in a higher proportion of sophisticated types (q) in the market. An advertising campaign aiming to improve consumer awareness regarding scenarios that their insurance plan does not cover falls in this category. The second category is policies that make early avoidance less costly (lower $\bar{a}$) by introducing competition in the secondary market. Such a policy would, for example, give consumers the option to borrow funds from a rival bank instead of using their bank's overdraft facility. The third category is policies that impose a cap on the fees that firms can extract from consumers in period 3 (lower $\bar{\bar{a}}$). A policy that places a limit on the fees that consumers can accumulate by being late on their credit card debt payments falls in this category.

At a first glance, increasing q would be expected to reduce firm profits by intensifying firm competition in fees. This is not the case at the maximum-differentiation local equilibrium; firms have no incentive to compete on fees, as competition on upfront prices is not exhausted. Higher q does not affect average travel costs either, as it does not affect firm locations. Intuitively, firms' incentive to locate further apart is driven by their desire to soften competition on upfront prices. The latter is independent of the proportion of sophisticated types in the market.

For a market at the intermediate-differentiation local equilibrium, increasing q has no effect on firm profits but affects average travel costs. Intuitively, this is because that equilibrium is characterised by the symmetric product locations with the smallest interfirm distance for which firms have no incentive to reduce add-on fees below the early-avoidance cost. Hence firms at the intermediate-differentiation local equilibrium respond to an increase in q by increasing interfirm distance sufficiently to maintain the same level of fee revenues. Whether higher q benefits consumers in terms of average travel costs depends on whether it shifts $\left(l_1^{int}, l_2^{int}\right)$ closer to or away from the location pair that minimises average travel costs, $\left(\frac{1}{4}, \frac{3}{4}\right)$. This is likely to put severe informational challenges to a regulator that is interested in how well products match consumer tastes, as expressing location pair $\left(\frac{1}{4}, \frac{3}{4}\right)$ in a real market terms may not be straightforward.

For a market at a minimum-differentiation local equilibrium increasing q has a negative effect on average firm profits. To see why, recall that average profits at this equilibrium are equal to the profits of a firm that earns $\bar{\bar{a}}$ from its naïve customers only. A lower proportion of naïve types, thus, reduces the population from which such fees can be extracted. An increase in q is neutral with respect to average travel costs, as offering perfect substitutes is the unique equilibrium when firms randomise across fees.

A policy intervention that lowers $\bar{a}$ has the direct effect of lowering the fees that consumers pay at the maximum-differentiation local equilibrium. However lowering fees has no effect on firm profits, as firms respond to such intervention by reducing the subsidy on upfront prices. Lower $\bar{a}$ has no effect on product differentiation either, as firms incentive to soften competition on upfront prices is independent of the fee revenues they are earning.

A policy that reduces $\bar{a}$ is effective at lowering firm profits when applied to a market in the intermediate-differentiation local equilibrium. This is because it reduces the fee revenues that firms retain due to the binding price floor. Lowering $\bar{a}$ also results in firms locating closer together, as the minimum interfirm distance for which firms achieve $a_i = \bar{a}$ decreases. Whether reducing interfirm distance benefits consumers in terms of average travel costs again depends on whether it shifts $\left(l_1^{int}, l_2^{int}\right)$ closer to or away from $\left(\frac{1}{4}, \frac{3}{4}\right)$.

Finally, lowering $\bar{a}$ has no effect on average profits at the minimum-differentiation local equilibrium, as average profits are driven by the late-avoidance cost. Accordingly, lowering $\bar{a}$ does not affect firms' incentive to offer products that are perfect substitutes.

Reducing the late-avoidance cost, $\bar{\bar{a}}$, does not constrain firm behaviour at the maximum-differentiation and the intermediate-differentiation local equilibria. Intuitively, this is because firms extract fees from both consumer types, so their behaviour is constrained by the value of the early-avoidance cost, $\bar{a}$. For a market at the minimum-differentiation local equilibrium, reducing $\bar{\bar{a}}$ strictly lowers firm profits, as it reduces the fee revenue that firms can extract from the naïve types if they choose the highest fee in the support. Firm's incentive to offer perfect substitutes remains, as average profit per customer is independent of locations.

4.4.2 Welfare comparison of local equilibria

Now suppose that, through market interventions, a regulator can steer the market towards one of the local equilibria presented in Proposition 2. A equilibrium-shifting intervention would, for instance, introduce enough competition in the secondary market that $\bar{a} > t$. This would eliminate the maximum-differentiation local equilibrium.

A regulator that is concerned solely with total surplus needs only compare the local equilibria in terms of average travel costs. It is, then, easy to see that the intermediate-differentiation local equilibrium is superior under the total surplus standard. This is since average travel costs are the highest when the market exhibits maximum or minimum product differentiation.

A regulator that is solely concerned with improving the position of the average consumer is interested in firm profits and average travel costs. Thus, in order to rank the local equilibria in terms of average consumer surplus we compare them in terms of average consumer expense. We define average consumer expense as the sum of industry profits and average travel costs.[68] Proposition 3 presents the findings of the comparison.

[68] The consumer expense that goes towards covering marginal costs remains the same in each local equilibrium, so we ignore it from the average consumer expense calculation. We also ignore the early-avoidance costs that sophisticated consumers incur at the minimum-differentiation local equilibrium if both firms set $a_i = \bar{\bar{a}}$, as that happens with almost zero probability.

Proposition 3. (Consumer surplus comparison).

(a) *The average consumer prefers the minimum-differentiation local equilibrium over the maximum-differentiation local equilibrium.*

(b) *The average consumer prefers the intermediate-differentiation local equilibrium over the maximum-differentiation local equilibrium.*

(c) *The average consumer prefers the intermediate-differentiation local equilibrium over the minimum-differentiation local equilibrium if* $\frac{\bar{a}q(t-\bar{a}q)}{4t} > \bar{a} - \bar{\bar{a}}(1-q)$.

Proof: *In appendix C.*

Comparing the minimum-differentiation equilibrium with the maximum-differentiation equilibrium is fairly straightforward, since both market structures entail the same average travel costs. The average consumer prefers a minimum-differentiation equilibrium, as that involves lower total prices overall. A comparison between the intermediate-differentiation equilibrium and either of the other two equilibria requires considering differences in both average total prices and average travel costs. In particular, the intermediate-differentiation equilibrium is preferred over the maximum-differentiation equilibrium in both aspects, as it entails lower average profits and non-extreme locations. Lastly, the intermediate-differentiation equilibrium is preferred over the minimum-differentiation equilibrium if the average consumer's gain in travel costs is higher than her detriment with respect to prices. The condition in part *(c)* of Proposition 3 reflects this. The right-hand side represents the net benefit to consumers in terms of travel cost in the intermediate-differentiation structure. The left-hand side represents the additional average surplus that consumers yield to firms in the intermediate-differentiation equilibrium.

Overall, Proposition 3 states that maximum differentiation is the worst outcome for consumers. This raises two crucial issues for regulators interested in consumer welfare. First, it shows that, rather than being a redeeming feature, dissipation of fee revenues may in fact indicate that consumer surplus is at its lowest. Consumers, on average, would be better off at an equilibrium that exhibits some retention of fee revenues, as such an equilibrium would be characterised by better deals due to the lower product differentiation. It follows that the markets most in need of monitoring may be the ones exhibiting large product differentiation, and upfront prices that are not constrained by a relevant price floor.

Second, it raises the possibility that the presence of naïve types may in fact be improving the deals available to all consumers by establishing a minimum-differentiation or an intermediate-differentiation equilibrium. Firm in a market populated solely by sophisticated types would instead offer maximally differentiated products at higher prices, behaving as in the standard Hotelling model with quadratic distance costs and endogenous prices.

Hence, Proposition 3 advises caution when intervening in markets with unanticipated fees. Inasmuch as policy interventions can shift the market to a different local equilibrium, regulating a market at, for instance, the minimum-differentiation local equilibrium can harm consumers and social surplus if that regulation establishes a maximum-differentiation equilibrium.

4.5 Conclusion

Despite economists' reliance on competition as the best means to protect consumers' interests, the ways in which firms' incentive to exploit myopic consumers affects and is affected by competition intensity remains understudied. Our paper attempts to shed some light on this area by studying how the firms' ability to retain profits from unanticipated add-on fees interacts with the degree of horizontal product differentiation in the market.

We show that the firms' locations on a Hotelling line critically influence price structure and the size of extra profits that firms can retain. High differentiation yields high prices, but results in full dissipation of add-on revenues. Intermediate differentiation permits retention of some extra profits, but results in lower prices overall. Low differentiation also permits retention of some extra profits, but can only establish a price equilibrium in mixed strategies. When firm locations are endogenous, firms' well-known incentive to soften competition by offering differentiated products conflicts with their incentive to maximise retained profits by offering closer substitutes. We show that this conflict can establish equilibria with maximum, minimum, or intermediate product differentiation.

Whether regulatory interventions are beneficial depends on the existing market equilibrium, and whether interventions can shift the market to a different equilibrium. Despite firms retaining extra profits, consumers are strictly better off in equilibria with lower differentiation, as they are associated with lower total prices and, in the case of intermediate-differentiation equilibrium, lower distance costs. The fact that these equilibria do not arise in a fully rational benchmark calls into question the established belief that market transparency improves market outcomes for consumers. On the contrary, it suggests that all consumers may be better off in markets with a proportion of myopic consumers.

Our model, while useful for addressing our research questions, does not come without limitations. The complexity of the pricing game with asymmetric locations does not permit us to explore the profitability of large deviations at this stage. We are leaving this problem for future work. Another interesting research question for future research would be to study how product differentiation incentives change when exploitation of consumer short-sightedness creates consumption distortions. A model attempting to answer this question would relax the assumptions of unit demand and full market coverage. Finally, our model provides a first look at how firms' product differentiation incentives change when consumers value add-on extras but mispredict their demand for them. Future work can look at this issue in greater depth by modelling a world in which consumers can choose to become sophisticated, or firms can choose to educate consumers about their likelihood of using add-on services.

4.6 Appendix C

Proof of Lemma 1.

First see that any $a_i > \bar{\bar{a}}$ is dominated by $a_i \leq \bar{\bar{a}}$. This is because every customer of firm i prefers to pay either $\bar{a}$ in period 1 or $\bar{\bar{a}}$ in period 2 over paying a_i in period 3. Therefore, firm i is better off with any $a_i < \bar{\bar{a}}$, as this guarantees at least some revenues from add-on fees.

Second see that any $a_i \in (\bar{a}, \bar{\bar{a}})$ is dominated by $a_i = \bar{\bar{a}}$. This is because naïve consumers pay any fee smaller or equal to $\bar{\bar{a}}$, whereas sophisticated consumers prefer to incur $\bar{a}$ in period 1 over paying $a_i \in (\bar{a}, \bar{\bar{a}}]$ in period 3.

Third, notice that for given f_i firm i's demand with $a_i = \bar{a}$ is the same as its demand with $a_i = \bar{\bar{a}}$. This is because sophisticated consumers that purchase from firm i with $a_i = \bar{\bar{a}}$ incur $\bar{a}$ to avoid the paying the add-on fee. Therefore, firm i's demand from sophisticates is described by $x_i^s(a_i^{es} = \bar{a})$, which is the same as firm i's demand from sophisticates with $a_i = \bar{a}$.

By assumption $\bar{a} > \bar{\bar{a}}(1-q)$, firm i prefers to earn $a_i = \bar{a}$ from both customer types over earning $a_i = \bar{\bar{a}}$ only from the naïve types. Thus, any price combination $(f_i, \bar{\bar{a}})$ is inferior to $(f_i, \bar{a})$. Therefore, there does not exist a pure strategy price equilibrium with $a_i > \bar{a}$.

Proof of Lemma 2.

Let $q < 1$. We will first show that in a pure strategy price equilibrium at least one constraint is always binding. To see this consider a pricing strategy with $f_i = \widehat{f_i} > c, a_i = \widehat{a_i} < \bar{a}$. Firm i's demand is given by $\mathrm{D}_i = \left(\frac{1}{2} - \frac{q(a_i - a^*)}{2t(2l_2-1)} - \frac{f_i - f^*}{2t(2l_2-1)}\right)$. Differentiation yields $\frac{\partial D_i}{\partial f_i} < \frac{\partial D_i}{\partial a_i}$, which suggests that demand's slope is more steep with respect to upfront prices that with respect to add-on fees. Then firm i can strictly increase its profits by playing $f_i = \widehat{f_i} - \epsilon, a_i = \widehat{a_i} + \epsilon$, maintaining the same mark-up while strictly increasing demand. Therefore, any pricing strategy with $f_i > c, a_i < \bar{a}$ is strictly dominated.

Since we know that at least one constraint is always binding, we can internalise one constraint each time and reduce the problem to a maximisation problem with one argument and one inequality constraint. First order conditions of the unconstrained maximisation problem are

$$f_i^* = \frac{t(2l_2 - 1) + (a^* - a_i)q + f^* - a_i + c}{2} \quad (2)$$

$$a_i^* = \frac{t(2l_2 - 1) + q(a^* - f_i + c) - f_i + f^*}{2q} \quad (3)$$

Suppose that the early-avoidance constraint is binding in equilibrium; that is (3) satisfies $\frac{t(2l_2-1)+q(a^*-f_i+c)-f_i+f^*}{2q} > \bar{a}$ $(3a)$. Then we can substitute $a^* = a_i = \bar{a}$ in (2) and rewrite it as $f_i^* = \frac{t(2l_2-1)+f^*-\bar{a}+c}{2}$ $(2a)$. Second order condition verifies that $(1a)$ describes a local maximum. Applying $f_i^* = f^*$ in $(2a)$ and rearranging yields $f^* = c - \bar{a} + t(2l_2 - 1)$. Comparing f^* with c gives that, given a binding early-avoidance constraint, the price floor constraint is binding if $l_2 < \frac{1}{2} + \frac{\bar{a}}{2t}$, and non-binding if $l_2 \geq \frac{1}{2} + \frac{\bar{a}}{2t}$. Put differently:

$$f^* = \begin{cases} c & if\ l_2 < \frac{1}{2} + \frac{\bar{a}}{2t} \\ c - \bar{a} + t(2l_2 - 1) & if\ l_2 \geq \frac{1}{2} + \frac{\bar{a}}{2t} \end{cases}$$

Substituting $f_i = f^* = c - \bar{a} + t(2l_2 - 1)$ back in $(3a)$ yields that the early-avoidance constraint is also binding. This is expected since at least one constraint is binding in equilibrium. Therefore, for $l_2 \geq \frac{1}{2} + \frac{\bar{a}}{2t}$ a pure strategy price equilibrium, if it exists, is characterised by $f^* = c - \bar{a} + t(2l_2 - 1)$ and $a^* = \bar{a}$. This is described in part *(a)* of Lemma 2. Substituting $f_i = f^* = c$ back in $(3a)$ and rearranging yields that the early-avoidance constraint $a^* = \bar{a}$ is binding if $l_2 > \frac{1}{2} + \frac{q\bar{a}}{2t}$. Therefore, for $\frac{1}{2} + \frac{\bar{a}}{2t} > l_2 > \frac{1}{2} + \frac{q\bar{a}}{2t}$ a pure strategy price equilibrium, if it exists, is characterised by $f^* = c$ and $a^* = \bar{a}$. This is described by part *(b)* of Lemma 2.

Repeating this process by initially assuming a binding price floor in equilibrium verifies that for $\frac{1}{2}+\frac{\bar{a}}{2t}>l_2>\frac{1}{2}+\frac{q\bar{a}}{2t}$ a pure strategy price equilibrium, if it exists, is characterised by $f^*=c$ and $a^*=\bar{a}$. It also shows that for $l_2<\frac{1}{2}+\frac{q\bar{a}}{2t}$ a pure strategy price equilibrium, if it exists, is characterised by $f^*= c$ and $a^*=\frac{t}{q}(2l_2-1)$. This is described by part *(c)* of Lemma 2.

Now let $q=1$. Then $\frac{\partial D_i}{\partial f_i}=\frac{\partial D_i}{\partial a_i}$, so demand is equally sensitive to both price components. Then we can rewrite the maximisation problem as $\max_{p_i}(p_i-c)\left(\frac{1}{2}-\frac{p_i-p^*}{2t(2l_2-1)}\right)$, where $p_i=f_i+a_i$. Since firms can allocate p_i in any way between f_i and a_i we can ignore the constraints. First and second order conditions yield that profits are maximised at $p_i^*=c+t(2l_2-1)$.

To find profits in each equilibrium substitute $f_i=f^*$ and $a_i=a^*$ in (1).■

Proof of Lemma 3.

Let $l_1=1-l_2$. Consider a candidate pure strategy price equilibrium with (c,a^*), where $a^*=\frac{t}{q}(2l_2-1)$. The proof proceeds in 5 steps. Step 1 establishes that a profitable deviation from (f^*,a^*) requires that firm i serves at least some naïves. Steps 2 and 3 establish that firm i maximises deviation profits by playing $(f_i,a_i)=(c,\bar{\bar{a}})$. Step 4 shows the l_2 for which such deviation is profitable. Step 5 shows that at a mixed strategy price equilibrium each firm i plays $f_i=c$ with probability 1 and earns $\Pi_i=\frac{\bar{a}(1-q)}{2}$.

Step 1*. A deviation in which firm i serves only sophisticated types is not profitable*.

A deviation by firm i with (f_i,a_i) such that all naïve types purchase from the rival requires that f_i is sufficiently greater that f_j. Then for firm i to attract any sophisticated types it must be that $a_i<a_j$. But such a price combination is suboptimal, since by Lemma 1, firm i is better off increasing a_i and reducing f_i by the same amount, increasing demand from naïve types and leaving demand from sophisticates unchanged.

Step 2. *For any f_i^d, firm i maximises deviation profits by choosing $a_i^d = \bar{\bar{a}}$.*

Let $a_{i\ ns}^d(f_i^d)$ represent the a_i^d above which the deviating firm enjoys no demand from sophisticates, and suppose $a_{i\ ns}^d(f_i^d) \in (a^*, \bar{a}]$. For deviations with $a_i^d < a_{i\ ns}^d$ the deviating firm enjoys some demand from sophisticates, so its profit maximisation problem is described by equation (3). Therefore, a pricing strategy with $a_i \neq a^*$ is strictly dominated by $a_i = a^*$. For deviations with $a_i^d \in [a_{i\ ns}^d(f_i^d), \bar{\bar{a}}]$ the deviating firm serves only naïve consumers. Then any $a_i^d \in [a_{i\ ns}^d(f_i^d), \bar{\bar{a}})$ is strictly dominated by $a_i^d = \bar{\bar{a}}$, as $\bar{\bar{a}}$ extracts the largest possible surplus from naïve consumers.

Now suppose $a_{i\ ns}^d(f_i^d) > \bar{a}$. Then a deviation with $a_i^d \leq \bar{a}$ cannot be profitable, as, again, the deviating firm's profit maximisation problem is described by equation (1). Therefore, the deviating firm i is better off playing $a_i = a^*$. For $a_i^d > \bar{a}$ any sophisticated customers of firm i avoid paying fees. This is because, from the sophisticated types' point of view, purchasing from firm i entails a future avoidance cost $\bar{a}$. Therefore, any $a_i^d \in (\bar{a}, \bar{\bar{a}})$ is strictly dominated by $\bar{\bar{a}}$, as $\bar{\bar{a}}$ extracts the largest possible surplus from naïve consumers.

Step 3. *For $a_i^d = \bar{\bar{a}}$, firm i maximises deviation profits by choosing $f_i^d = c$.*

Consider a price combination $(f_i^d, \bar{\bar{a}})$ such that both types purchase from firm i. Firm i's deviation profits are given by $(f_i^d + \bar{\bar{a}} - c)(1-q)\left(\frac{1}{2} - \frac{f_i^d - c}{2t(2l_2-1)}\right) + (f^d - c)q\left(\frac{1}{2} - \frac{f_i^d + \bar{a} - c - a^*}{2t(2l_2-1)}\right)$. Maximisation with respect to f_i^d yields $f_i^{d^*} = \frac{t(2l_2-1)+c+\bar{\bar{a}}q-\bar{\bar{a}}+c-\bar{a}q+a^*q}{2}$. Since $f_i^{d^*}$ is increasing in l_2. Then, comparing $f_i^{d^*}$ evaluated at $l_2 = \frac{1}{2} + \frac{q\bar{a}}{2t}$, with c gives that $f_i^{d^*} \leq c$ for all $l_2 \in l_2^L$, $a \in [a^*, \bar{a}]$ if $\bar{\bar{a}}(1-q) \geq \bar{a}q$.

Consider a deviation at $f_i^d, \bar{\bar{a}}$ such that only naïve types purchase from firm i. Firm i's deviation profit with $a_i = \bar{\bar{a}}$ is $(f_i^d + \bar{\bar{a}} - c)(1-q)\left(\frac{1}{2} - \frac{f_i^d - c}{2t(2l_2-1)}\right)$. Maximisation with respect to f_i^d yields $f_i^{d^{**}} = \frac{t(2l_2-1)+c-\bar{\bar{a}}+c}{2}$. Then, comparing $f_i^{d^{**}}$ evaluated at

$l_2 = \frac{1}{2} + \frac{q\bar{a}}{2t}$, with c gives that $f_i^{d^{**}} \leq c$ for all $l_2 \leq \frac{1}{2} + \frac{\bar{\bar{a}}}{2t}$. But $\frac{1}{2} + \frac{\bar{\bar{a}}}{2t}$ lies to the right of l_2^L. Therefore, for all $l_2 \in l_2^L$, $f_i^{d^{**}} = c$.

Step 4. *Deviation with* $(f_i, a_i) = (c, \bar{\bar{a}})$ *is profitable if and only if* $l_2 \leq \frac{1}{2} + \frac{\bar{\bar{a}}q(1-q)}{2t}$.

At $(c, \bar{\bar{a}})$, the deviating firm's profits are $\frac{\bar{\bar{a}}(1-q)}{2}$. Then, the deviation is profitable if $\frac{\bar{\bar{a}}(1-q)}{2} > \frac{a_L^*}{2}$. Expressing the inequality in terms of l_2 gives that the deviation is profitable if $l_2 \leq \frac{1}{2} + \frac{\bar{\bar{a}}q(1-q)}{2t}$. Therefore, there exists a pure strategy equilibrium with prices as described in Lemma 2 if $l_2 > \frac{1}{2} + \frac{\bar{\bar{a}}q(1-q)}{2t}$.

Step 5. *There exists a mixed strategy price equilibrium in which firms play* $f_i = c$ *with probability 1 and earn, on average,* $\Pi_i = \frac{\bar{\bar{a}}(1-q)}{2}$.

To show that $f_i = c$ is part of an equilibrium strategy we only need to show that firm i has no incentive to play $f_i > c$ if firm j plays $f_j = c$. Recall that firm i can only play $f_i > c$ if $a_i \in \{\bar{a}, \bar{\bar{a}}\}$. Otherwise, firm i can strictly increase its profit by marginally reducing the upfront price and increasing the add-on fee by the same value. Therefore, we only need to consider the profit maximising value of f_i given $a_i = \bar{a}$ and $a_i = \bar{\bar{a}}$.

First, suppose $a_i = \bar{\bar{a}}$ and $f_j = c$. Then replacing a^* with $a_j^{es} \in [a^*, \bar{a}]$ and repeating step 3 gives that $f_i = c$ maximises profit for any a_j if $a_i = \bar{\bar{a}}$, $f_j = c$. To see why a_j^{es} is bound above by $\bar{a}$ recall that sophisticated types prefer to incur $\bar{a}$ and avoid fees over paying $a_j > \bar{a}$.

Next, suppose $a_i = \bar{a}$ and $f_j = c$, and suppose that both types purchase from firm i. Then firm i's profits are given by $(f_i + \bar{a} - c)\left(\frac{1}{2} - \frac{f_i - \underline{f} + q\left(\bar{a} - a_j^{es}\right)}{2t(2l_2 - 1)}\right)$. Maximisation with respect to f_i yields $f_i^* = \frac{t(2l_2-1)+c-\bar{a}+c-\bar{a}q+a_j^{es}q}{2}$. Notice that f_i^* is increasing in l_2, a_j^{es}. Therefore, replacing l_2 with the upper bound of $l_2^L, \frac{1}{2} + \frac{q\bar{a}}{2t}$ and a_j^{es} by its highest value, $\bar{a}$, gives the highest value of f_i^*. Comparing f_i^* evaluated at $l_2 = \frac{1}{2} + \frac{q\bar{a}}{2t}$, $a_j^{es} = \bar{a}$ with c gives that $f_i^{d^*} \leq c$.

Now suppose that only naïve types purchase from firm i. Firm i's deviation profit with $a_i = \overline{a}$ is $(f_i + \overline{a} - c)(1-q)\left(\frac{1}{2} - \frac{f_i - c}{2t(2l_2 - 1)}\right)$. Maximisation with respect to f_i yields $f_i^{**} = \frac{t(2l_2-1)+\underline{f}-\bar{a}+c}{2}$. Comparison with c yields that $f_i^{**} \leq c$ for all $l_2 \leq \frac{1}{2} + \frac{\overline{a}}{2t}$. But $\frac{1}{2} + \frac{\overline{a}}{2t}$ lies to the right of l_2^L. Hence $f_i^{**} = c$ for all $l_2 \in l_2^L$. Therefore, $f_i = c$ maximises profit for any a_j if $a_i = \overline{a}$, $f_j = c$.

By step 2 we have that $\bar{\bar{a}}$ is always an element of the fee support in the mixed strategy price equilibrium. Then since profits are identical at every fee in the support, firm i's profits at a mixed strategy price equilibrium are $\frac{\bar{\bar{a}}(1-q)}{2}$.■

Proof of Proposition 2.

Let $l_2 \in l_2^H$. In order to calculate the marginal relocation incentives of each firm for a given symmetric location pair we need to express equilibrium profits in terms of l_1, l_2. To do so we will solve the maximisation problem of section 4.3.1.1 without making use of the symmetry condition $l_1 = 1 - l_2$. Recall that firm demands at a symmetric pure strategy price equilibrium are given by $D_1 = qx^s + (1-q)x^n$, and $D_2 = q(1-x^s) + (1-q)(1-x^n)$. Thus firm 1 maximises $\Pi_1 = (f_1 + a_1 - c)\left(\frac{l_1+l_2}{2} - \frac{f_1-f_2}{2t(l_2-l_1)} - \frac{q(a_1-a_2)}{2t(l_2-l_1)}\right)$; firm 2 maximises $\Pi_2 = (f_2 + a_2 - c)\left(1 - \frac{l_1+l_2}{2} + \frac{f_1-f_2}{2t(l_2-l_1)} + \frac{q(a_1-a_2)}{2t(l_2-l_1)}\right)$

Let $l_2 \in l_2^H$. By Proposition 1 we know that each firm sets $a_i = \bar{a}$ and $f_i < c$. Thus equilibrium prices are $f_{H1}^* = argmax(f_1 + \bar{a} - c)\left(\frac{l_1+l_2}{2} - \frac{f_1-f_2}{2t(l_2-l_1)}\right)$ and $f_{H2}^* = argmax(f_2 + \bar{a} - c)\left(1 - \frac{l_1+l_2}{2} + \frac{f_1-f_2}{2t(l_2-l_1)}\right)$. Solving simultaneously yields $f_{H1}^* = c - \bar{a} + \frac{t(l_2-l_1)(2+l_1+l_2)}{3}$ and $f_{H2}^* = c - \bar{a} + \frac{t(l_2-l_1)(4-l_1-l_2)}{3}$. Substituting back to each firm's profit function yields $\Pi_{H1}^* = \frac{t(l_2-l_1)(2+l_1+l_2)^2}{18}$ and $\Pi_{H2}^* = \frac{t(l_2-l_1)(4-l_1-l_2)^2}{18}$. Differentiation yields $\frac{d\Pi_{H1}^*}{dl_1} < 0$ and $\frac{d\Pi_{H2}^*}{dl_2} > 0$; that is, each firms prefers to locate further away from its rival. Now notice that (l_1^{max}, l_2^{max}) satisfies $l_2 \in l_2^H$. Since (l_1^{max}, l_2^{max}) describes locations at the ends of the line we can conclude that any local deviation from (l_1^{max}, l_2^{max}) is unprofitable, as firms cannot locate any further apart. Replacing l_1^{max} and l_2^{max} in equilibrium prices and profits gives the desired f_{max}, a_{max}, and Π_{max}.

Now let $l_2 \in l_2^I$. By Proposition 1 we know that each firm sets $a_i = \bar{a}$ and $f_1 = c$. Profits of firm 1 are given by $\Pi_{I1}^* = \frac{\bar{a}(l_1+l_2)}{2}$ and $\Pi_{I2}^* = \frac{\bar{a}(2-l_1-l_2)}{2}$. Differentiation yields $\frac{d\Pi_{I1}^*}{dl_1} > 0$ and $\frac{d\Pi_{I2}^*}{dl_2} < 0$; that is, each firm prefers to offer a product that is a closer substitute to its rival's product.

Next let $l_2 \in l_2^L$. By Proposition 1 we know that each firm sets $f_1 = c$. Then $a_{L1}^* = argmax(a_1)\left(\frac{l_1+l_2}{2} - \frac{q(a_1-a_2)}{2t(l_2-l_1)}\right)$ and $a_{L2}^* = argmax(a_2)\left(1 - \frac{l_1+l_2}{2} + \frac{q(a_1-a_2)}{2t(l_2-l_1)}\right)$. Solving simultaneously yields $a_{L1}^* = \frac{t(l_2-l_1)(2+l_1+l_2)}{3q}$ and $a_{L2}^* = \frac{t(l_2-l_1)(4-l_1-l_2)}{3q}$. Substituting back to each firms profit functions yields $\Pi_{L1}^* = \frac{t(l_2-l_1)(2+l_1+l_2)^2}{18q}$ and $\Pi_{L2}^* = \frac{t(l_2-l_1)(4-l_1-l_2)^2}{18q}$. Differentiation yields $\frac{d\Pi_{L1}^*}{dl_1} < 0$ and $\frac{d\Pi_{L2}^*}{dl_2} > 0$; that is, each firm prefers to locate further away from its rival.

To see why location pair $\left(l_1^{int}, l_2^{int}\right)$ constitutes a local equilibrium notice that l_2^{int} lies on the bound between l_2^I and l_2^L. Therefore, evaluating firm 2's profits at l_2^{int} from the right yields $\Pi_2|_{l_2=l_2^{int+}} = \Pi_{I2}^* = \frac{\bar{a}(2-l_1-l_2)}{2}$, while evaluating them from the left yields $\Pi_2|_{l_2=l_2^{int-}} = \Pi_{L2}^* = \frac{t(l_2-l_1)(4-l_1-l_2)^2}{18q}$. It follows that the right-side derivative with respect to l_2 is $\frac{d\Pi_2^*}{dl_2}\Big|_{l_2=l_2^{int+}} = \frac{d\Pi_{I2}^*}{dl_2} < 0$, while the left-side derivative with respect to l_2 is $\frac{d\Pi_2^*}{dl_2}\Big|_{l_1=l_2^{int-}} = \frac{d\Pi_{L2}^*}{dl_2} > 0$. The same holds for firm 1 by symmetry. Thus we can conclude that location pair $\left(l_1^{int}, l_2^{int}\right)$ constitutes a local equilibrium.

Finally, let $l_2 \in l_2^M$. As established in Proposition 1, in the price equilibrium firm i's profits are on average equal to its profits if it sets $\bar{\bar{a}}$ and serves all naïve types at its turf and half the naïve types located between firm i and firm j. Then we can express firm 1's expected profits by $\Pi_{M1}^* = \frac{\bar{\bar{a}}(1-q)(l_1+l_2)}{2}$ and firm 2's expected profits by

$\Pi^*_{M2} = \frac{\bar{\bar{a}}(1-q)(2-l_1-l_2)}{2}$. Differentiation yields $\frac{d\Pi^*_{M1}}{dl_1} > 0$ and $\frac{d\Pi^*_{M2}}{dl_2} < 0$; firms prefer to offer more homogeneous products. Now notice that (l_1^{min}, l_2^{min}) satisfies $l_2 \in l_2^M$. It, then, follows from Proposition 1 that firms set $f_{min} = c$ and choose a_{min} randomly. Since (l_1^{min}, l_2^{min}) describes locations at the centre of the line we can conclude that any local deviation from (l_1^{min}, l_2^{min}) is unprofitable, as firms cannot locate any closer together. Therefore, (l_1^{min}, l_2^{min}) constitutes a local equilibrium. Replacing l_1^{min} and l_2^{min} in Π^*_{M1} and Π^*_{M2} gives the desired Π_{min}.

The proof for the characterisation of the support of a_{min} is as follows. Let a^- denote the lower bound in the support. If firm i sets $a_i = a^-$ it serves half the naïve types in the market and every sophisticated type. Its profits are given by $\Pi_i(a^-) = a^-\left(\frac{(1-q)}{2} + q\right)$. Since firm i must be indifferent between any fee in the support, a^- satisfies $\Pi_i(a^-) = \Pi_{min}$. Substituting and rearranging gives $a^- = \bar{\bar{a}}\frac{(1-q)}{1+q}$. To see why there exists a gap in the fee support notice that firm i has no incentive to play $a_i \in [\bar{a}, \bar{\bar{a}})$. For $a_i \in (\bar{a}, \bar{\bar{a}})$ any sophisticated types that purchase from firm i avoid fees. Therefore, firm i can increase its profits by setting $a_i = \bar{\bar{a}}$, extracting higher fee revenues from naïve types while maintaining the same demand from sophisticated types. For $a_i = \bar{a}$ firm j has a lower fee with probability almost equal to 1, so it almost surely attracts all sophisticated types. Then firm i is again better off playing $a_i = \bar{\bar{a}}$, extracting higher fee revenues from any naïve types it serves while serving sophisticated types with almost zero probability.■

Proof of Proposition 3.

Recall that χ represents consumer locations. Average consumer expense at the maximum-differentiation market structure is given by $K_{max} = 2\Pi_{max} + t\left(\int_0^{\frac{1}{2}}(0-\chi)^2 d\chi + \int_{\frac{1}{2}}^1 (1-\chi)^2 d\chi\right) = t + \frac{t}{12}$. Average consumer expense at the minimum-differentiation market structure is given by $K_{min} = 2\Pi_{min} + t\int_0^1 \left(\frac{1}{2} - \chi\right)^2 d\chi = \bar{\bar{a}}(1-q) + \frac{t}{12}$. Average consumer expense at the intermediate-differentiation market structure is given by $K_{int} = 2\Pi_{int} + t\left(\int_0^{\frac{1}{2}}\left(\frac{1}{2} - \frac{q\bar{a}}{2t} - \chi\right)^2 d\chi + \int_{\frac{1}{2}}^{1}\left(\frac{1}{2} + \frac{q\bar{a}}{2t} - \chi\right)^2 d\chi\right) = \bar{a} + \frac{t^2 - 3\bar{a}qt + 3\bar{a}^2q^2}{12t}$. Comparison yields the desired results. ■

References

Agarwal, S., Chomsisengphet, S., Mahoney, N. and Stroebel, J. (2015). "Regulating consumer financial products: Evidence from credit cards", *Quarterly Journal of Economics*, 130(1): 111-164.

Agarwal, S., Chomsisengphet, S., Mahoney, N. and Stroebel, J. (2014). "A simple framework for estimating consumer benefits from regulating hidden fees", *Journal of Legal Studies*, 43(2): 239-252.

Ahmetoglu, G., Furnham, A. and Fagan, P. (2014). "Pricing practices: A critical review of their effects on consumer perceptions and behaviour", *Journal of Retailing and Consumer Services*, 21: 696-707.

Armstrong, M. (2015). "Search and ripoff externalities", *Review of Industrial Organisation*, 47: 273-302.

Armstrong, M. and Chen, Y. (2009). "Inattentive consumers and product quality", *Journal of European Economic Association*, 7(2-3): 411-422.

Armstrong, M. and Vickers, J. (2012). "Consumer protection and contingent charges", *Journal of Economic Literature*, 50(2): 477-493.

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.

Bordalo, P., Gennaioli, N. and Shleifer, A. (2016). "Competition for attention", *Review of Economic Studies*, 83: 481-513.

Borenstein, S., MacKie-Mason, J. and Netz, J. (2000). "Exercising market power in proprietary aftermarkets", *Journal of Economics and Management Strategy*, 9(2): 157-188.

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.

Carlton, D. W. and Waldman, M. (2002). "The strategic use of tying to preserve and create market power in evolving industries", *RAND Journal of Economics*, 33(2): 194-220.

Carlton, D. W. and Waldman, M. (2009). "Competition, monopoly, and aftermarkets", *Journal of Law, Economics and Organization*, 26(1): 54-91.

Chen, Y. (1997). "Equilibrium product bundling", *Journal of Business*, 70: 85-103.

Chioveanu, I. and Zhou, J. (2013). "Price competition with consumer confusion", *Management Science*, 59(11): 2450-2469.

Cui, Y., Duenyas, I., and Sahin, O. (2018). "Unbundling of ancillary service: how does price discrimination of main service matter", *Manufacturing and Service Operations Management*, 20(3): 455-466.

d'Aspremont, C., Gabszewicz, J. J., and Thisse, J.-F. (1979). "On Hotelling's 'Stability in competition'", *Econometrica*, 47: 1145-1150.

Dahremoller, C. (2013). "Unshrouding for competitive advantage", *Journal of Economics and Management Strategy*, 22(3): 551-568.

de Meza, D. and Reyniers, D. (2012). "Every shroud has a silver lining: The visible benefits of hidden surchargers", *Economics Letters*, 116(2): 151-153.

DellaVigna, S. and Malmendier, U. (2004). "Contract design and self-control: theory and evidence", *Quarterly Journal of Economics*, 119(2): 353–402.

Dixit, A. (1986). "Comparative statics for oligopoly", *International Economic Review*, 27(1): 107-122.

Ellison, G. (2005). “A model of add-on pricing”, *Quarterly Journal of Economics*, 120(2): 585-637.

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”, *RAND Journal of Economics*, 43(3): 417-441.

Farrell, J. and Klemperer, P. (2007). “Coordination and lock-in: Competition with switching costs and network effects”. In Armstrong, M. and Porter, R. (eds.), *Handbook of Industrial Organisation,* 3: 1967-2072.

Federal Communications Commission (2010). *White paper on bill shock*. Consumer and Governmental Affairs Bureau.

Financial Conduct Authority (2015). “Consumer vulnerability”, Occasional Paper No. 8.

Fletcher, A. (2016). *The role of demand-side remedies in driving effective competition – a review for which?* at https://www.staticwhich.co.uk/documents/pdf/the-role-of-demand-side-remedies-in-driving-effective-competition-456067.pdf (accessed 16 July 2018).

Fruchter, G., Gerstner, E. and Dobson, P. (2011). “Fee or free? How much to add on for an add-on?”, *Marketing Letters,* 22(1): 65-78.

Gabaix, X. and Laibson, D. (2006). “Shrouded attributes, consumer myopia and information suppression in competitive markets”, *Quarterly Journal of Economics,* 121(2): 505–540.

Gamp, T. and Krahmer, D. (2018). “Deception and competition in search markets”, working paper.

Gathergood, J., Sakaguchi, H., Stewart, N. and Weber, J. (2018). “How do consumers avoid penalty fees? Evidence from credit cards”, working paper.

Geng, X. and Shulman, J. (2015). “How costs and heterogeneous consumer price sensitivity interact with add-on pricing”, *Production and Operations Management*, 24(12): 1870-1882.

Geng, X., Tan, Y. and Wei, L. (2018). "How add-on pricing interacts with distribution contracts", *Production and Operations Management*, 27(4): 605-623.

Gomes, R. and Tirole, J. (2018). "Missed sales and the pricing of ancillary goods", Forthcoming in *Quarterly Journal of Economics*.

Greenleaf, E.A., Johnson, E.J., Morwitz, V.G. and Shalev, E. (2016). "The price does not include additional taxes, fees, and surcharges: A review of research on partitioned pricing", *Journal of Consumer Psychology*, 26(1): 105-124.

Grossman, G. M. and Shapiro, C. (1984). "Informative advertising with differentiated products", *Review of Economic Studies*, 51: 63-81.

Grubb, M. D. (2015a). "Behavioral consumers in industrial organization: an overview", *Review of Industrial Organisation*, 47(3): 247-258.

Grubb, M. D. (2015b). "Consumer inattention and bill-shock regulation", *Review of Economic Studies*, 82(1): 219–257.

Grubb, M. D. and Osborne, M. (2015). "Cellular service demand: Biased beliefs, learning, and bill shock", *American Economic Review*, 105(1): 234-271.

Heidhues, P. and Koszegi, B. (2010). "Exploiting naïvete about self-control in the credit market", *American Economic Review*, 100(5): 2279-2303.

Heidhues, P. and Koszegi, B. (2015). "On the welfare costs of naïvete in the us credit-card market", *Review of Industrial Organisation*, 47(3): 341-354.

Heidhues, P. and Koszegi, B. (2017). "Naïvete-based price discrimination", *Quarterly Journal of Economics*, 132(2): 1019-1054.

Heidhues, P. and Koszegi, B. (2018). "Behavioral industrial organization", working paper.

Heidhues, P., Johnen, J. and Koszegi, B. (2018). "Browsing versus studying: A pro-market case for regulation", working paper.

Heidhues, P., Koszegi, B. and Murooka, T. (2016). "Exploitative innovation", *American Economic Journal: Microeconomics*, 8(1): 1-23.

Heidhues, P., Koszegi, B. and Murooka, T. (2017). "Inferior products and profitable deception", *Review of Economic Studies*, 84(1): 323-356.

Hotelling, H. (1929). “Stability in competition”, *Economic Journal*, 39(153): 41-57.

Klemperer, P. (1987). “Markets with consumer switching costs”, *Quarterly Journal of Economics*, 102(2): 375-394.

Ko, J. and Williams, J. (2017). “The effects of regulating hidden add-on costs”, *Journal of Money, Credit and Banking*, 49(1): 39-74.

Kosfeld, M. and Schuwer, U. (2017). “Add-on pricing in retail financial markets and the fallacies of consumer education”, *Review of Finance,* 21(3): 1189-1216.

Kühn, K.-U. and van Reenen, J. (2009). “Interoperability and market foreclosure in the European Microsoft case”. In Lyons, B. (ed.), *Cases in European Competition Policy: The Economic Analysis.* Cambridge: Cambridge University Press.

Lal, R. and Matutes, C. (1994). “Retail pricing and advertising strategies”, *Journal of Business*, 67(3): 345-370.

Lin, S. (2017). “Add-on policies under vertical differentiation: Why do luxury hotels charge for internet while economy hotels do not?”, *Marketing Science*, 36(4): 610-625.

Miao, C.-H. (2010). “Consumer myopia, standardization and aftermarket monopolization”, *European Economic Review*, 54(7): 931-946.

Michel. C. (2017). “Market regulation of voluntary add-on contracts”, *International Journal of Industrial Organization*, 54: 239-268.

Michel. C. (2018). “Contractual structures and consumer misperceptions”, *Journal of Economics and Management Strategy*, 27(2): 188-205.

Murooka, T. (2015). “Deception under competitive intermediation”, working paper.

Neven, D. (1985). “Two stage (perfect) equilibrium in Hotelling’s model”, *Journal of Industrial Economics*, 33(3): 317-325.

Office of Fair Trading (2008). *Personal current accounts in the UK: an OFT market study*.

Office of Fair Trading (2013). *Partitioned pricing research: The economics of partitioned pricing*.

Piccione, M. and Spiegler, R. (2012). “Price competition under limited comparability”, *Quarterly Journal of Economics*, 127(1): 97-135.

Schultz, C. (2004). “Market transparency and product differentiation”, *Economics Letters*, 83: 173-178.

Schultz, C. (2005). “Transparency on the consumer side and tacit collusion”, *European Economic Review, 49: 279-297*.

Shapiro, C. (1995). “Aftermarkets and consumer welfare: Making sense of Kodak”, *Antitrust Law Journal*, 63(2): 483-511.

Shelanski, H.A., Farrell, J., Hanner, D., Metcalf, C.J., Sullivan, M.W., Wendling, B.W. (2012). “Economics at the FTC: Drug and PBM mergers and drip pricing”, *Review of Industrial Organization*, 41: 303-319.

Spiegler, R. (2006). “The market for quacks”, *Review of Economic Studies*, 73: 1113-1131.

Spiegler, R. (2014). “Competitive framing”, *American Economic Journal: Microeconomics*, 6(3): 35-58.

Spiegler, R. (2016). “Choice complexity and market competition”, *Annual Review of Economics*, 8: 1-25.

Stole, L. (2007). “Price discrimination and imperfect competition”. In Armstrong, M. and Porter, R. (eds.), *Handbook of Industrial Organisation*, 3: 2221-2299.

Tremblay, V. J. and Polasky, S. (2002). “Advertising with subjective horizontal and vertical product differentiation”, *Review of Industrial Organization*, 20: 253-265.

Varian, H. (1980). “A model of sales”, *American Economic Review*, 70: 651-659.

Verboven, F. (1999). “Product line rivalry and market segmentation – with an application to automobile optional engine pricing”, *Journal of Industrial Economics*, 47(4): 399-425.

Wenzel, T. (2014). “Consumer myopia, competition and the incentives to unshroud add-on information”, *Journal of Economic Behavior and Organization*, 98: 89-96.

Wenzel, T. (2015). “Partial shrouding in asymmetric markets”, *Economics Letters*, 136: 223-226.

Weyl, G. E. and Fabinger, M. (2013). “Pass-through as an economic tool: Principles of incidence under imperfect competition”, *Journal of Political Economy,* 121(3): 528-583.

Whinston, M. D. (1990). “Tying, foreclosure, and exclusion”, *American Economic Review,* 80(4): 837–59.

Wilson, C. M. (2010). “Ordered search and equilibrium obfuscation”, *International Journal of Industrial Organisation,* 28(5): 496-506.

Wolinsky, A. (1983). “Prices as signals of product quality”, *Review of Economic Studies,* 50(4): 647-658.

Zegners, D. and Kretschmer, T. (2017). “Competition with aftermarket power when consumers are heterogeneous”, *Journal of Economics and Management Strategy,* 26(1): 96-122.

Zenger, H. (2013). “Why firms’ exploitation of consumer myopia may benefit myopic consumers”, *Economics Letters,* 118: 307-309.