Product Lines and Price Discrimination in Markets with Information Frictions

October 1, 2020

Abstract

A well-known principle in economics is that firms differentiate their product offerings in order to relax competition. However, in this paper we show that information frictions can invalidate this principle. We build a duopolistic competition model of second-degree price discrimination with information frictions in which (i) an equilibrium always exists with overlapping product qualities, whereas (ii) an equilibrium with non-overlapping product qualities exists only if information frictions and the costs of providing high quality are sufficiently small. As a consequence, reasons other than an attempt to soften competition should explain why firms in some cases carry non-overlapping product lines.

Keywords: product strategy, pricing strategy, second-degree price discrimination, search, vertical differentiation, retail competition.

1 Introduction

Different reasons have been advanced to explain why competing firms carry overlapping product lines in some markets, but not others. Pervasive among reasons for the latter situation is the well-known Chamberlinian principle that firms seek to differentiate their products in order to relax competition (Chamberlin, 1933). Champsaur and Rochet (1989) (CR, hereafter) formalized this principle in a model in which quality choices are followed by price competition.[1] CR showed that an equilibrium exists in which

[1]Shaked and Sutton (1982) and Moorthy (1988) formalized the same idea in a model similar to Champsaur and Rochet's (1989), with the difference that firms are allowed to offer one quality only. Thus, in Shaked and Sutton (1982) and Moorthy (1988), discrimination between consumers at the firm level is not possible.

firms choose non-overlapping qualities because the incentives to soften price competition dominate the incentives to better discriminate consumers with heterogeneous preferences for quality. They also argued that, not surprisingly, this differentiation principle should weaken as competition declines.

The general goal of this paper is to understand whether the Chamberlinian principle of product differentiation is robust to introducing imperfect competition linked to information or search frictions. For this purpose, and in line with previous papers, we assume that consumers' preferences and product qualities are such that a monopolist would find it optimal to discriminate consumers by carrying the full product line. In such a set-up, would competition in markets with information frictions still give rise to non-overlapping product choices, as it does in the case of frictionless markets?

We show that the Chamberlinian principle survives in markets with information frictions (i.e., an equilibrium exists in which firms carry products with non-overlapping qualities), but only as long as such frictions and the costs of providing high quality are small enough. In contrast, the equilibrium in which the two competing firms carry full product lines, and thus compete head-to-head, always exists. Since no market is immune to information frictions, our finding provides an important lesson for applied work: competing firms may carry non-overlapping product lines, but likely for reasons other than the Chamberlinian incentive to soften competition. Hence, researchers should rely on other models to account for asymmetries in product lines.

Consumers who are imperfectly informed about firms' product prices and qualities cannot choose their best option without incurring search costs to learn about and compare all options. Since the seminal work of Diamond (1971), the search literature has shown that introducing information frictions can substantially affect competition. However, unlike CR, this literature has mostly neglected the possibility that firms engage in price discrimination through their quality choices.[2] The main goal of this paper is to understand the interaction between these two forces: competition and price discrimination.

By introducing information frictions *à la* Varian (1980) (i.e., a fraction of consumers are *uninformed* about firms' product offerings and prices), we show that, if the costs of providing high quality are large enough, an arbitrarily small number of uninformed consumers is all it takes to rule out an equilibrium in which firms offer non-overlapping product choices.[3] Intuitively, the presence of uninformed consumers induces the firm

[2]Unlike the current paper, in which we model second-degree price discrimination, Fabra and Reguant (2020) allow for third-degree price discrimination in markets with search costs.

[3]The online appendix shows the same results if we endogenize the fraction of uninformed consumers by explicitly accounting for search costs. For this purpose, we follow the fixed-sample approach of Burdett and Judd (1983) but allow consumers to differ in terms of their search costs.

carrying low quality products to deviate by also carrying high quality products in order to better discriminate consumers without fear of sacrificing profits on low quality items. If providing high quality is not too costly, a sufficiently large number of uninformed consumers also rule out the equilibrium with non-overlapping qualities as the gains from price discrimination outweigh the gains from softening competition. Instead, the equilibrium in which firms offer overlapping qualities always exists, no matter whether there are none, few, or many uninformed consumers,[4] and no matter how costly it is to provide high quality. In this sense, the equilibrium with overlapping qualities is particularly robust (and for a large set of parameter values, unique),[5] whereas the equilibrium with non-overlapping qualities proposed by CR is not.[6] Another compelling reason for focusing on the equilibrium with overlapping qualities is that it naturally converges to the Bertrand equilibrium as information frictions vanish.[7]

Beyond investigating the effects of information frictions on firms' quality choices, we also seek to understand the effects of information frictions on equilibrium pricing in general. In this sense, we extend Varian (1980)'s equilibrium to a multi-product firm setting. In particular, we show that the incentive compatibility constraints faced by multi-product firms introduce an important departure from Varian (1980): the prices of various goods sold in a store cannot be independent of each other. This has several

[4]The equilibrium with overlapping qualities trivially exists in the absence of information frictions, in which case Bertrand pricing prevails. CR focus only on equilibria with strictly positive profits, and hence abstract from this equilibrium. One of our contributions is to show that the overlapping equilibrium always exists while the non-overlapping equilibrium may disappear altogether, even when the fraction of uninformed consumers is arbitrarily small.

[5]If we introduced search costs *à la* Diamond, in which consumers search firms sequentially at some positive cost, firms would also carry overlapping product lines. However, because of the Diamond paradox, consumers would not search and firms would not compete among themselves. Therefore, this assumption would not be appropriate to analyze the interaction between competition and price discrimination. Varian's approach avoids this paradox, giving rise to comparative statics that replicate empirical findings regarding search behavior and price patterns.

[6]One important difference between our model and CR's is that we consider a model with two qualities while they allow firms to choose a quality range from a continuum of options. This difference does not drive the difference in results. The online appendix shows that as soon as we introduce an arbitrarily small number of uninformed consumers into CR's model, the low-quality firm has incentives to deviate outside the quality gap by carrying high-quality items that overlap with items offered by the high-quality firm. Given this deviation, how to differentiate the high from the low-quality firm is no longer obvious. In contrast, as also shown in the online appendix, introducing a small amount of horizontal differentiation into CR's model shrinks the quality gap just a bit, and the low-quality firm remains as such.

[7]As suggested by a referee, yet another reason for focusing on the equilibrium with full overlap under information frictions (no matter how small) is that the presence of more than two firms makes it the unique equilibrium, even if the costs of providing high quality are not large enough.

implications for pricing behavior. For instance, if both firms carry high and low quality goods and competition intensifies sufficiently (which Varian refers to as periods of *sales*), the firms will reduce the price of the high-quality good relative to that of the low-quality good to the extent that the incentive compatibility constraint no longer binds. The reason is that firms' incentives to compete for consumers with a preference for high quality may dominate their incentives to minimize these types' information rents. Additionally, incentive compatibility considerations imply that multi-product firms tend to charge lower prices, on average, compared to single-product firms, contrary to prices charged by single-product versus multi-product monopolists under complete information.

Related Literature Our paper relates to two strands of the literature: papers that analyze competition with search frictions, and papers that characterize quality choices under imperfect competition.[8] The vast majority of recent search models allow for horizontal product differentiation (e.g., Wolinsky, 1985; Anderson and Renault, 1999; Kuksov, 2004; and Bar-Isaac *et al.* 2012).[9] Furthermore, they assume either that each firm will carry a single product or, if they allow firms to carry several products, that consumers will search for more than one product ('multi-product search').[10] In these models, consumers differ in their preference for buying all goods in the same store ('one-stop shopping') rather than in their quality preferences.[11]

With regard to our model, these differences are relevant. In the first type of search models, the single-product firm assumption leaves no scope for price discrimination within the firm. Hence, pricing is solely driven by competitive forces. A notable example is given by Kuksov (2004), who develops a Hotelling model with search costs in which two single-product firms choose their locations along a line. He finds that firms increasingly differentiate their products as reduced search costs intensify competition, potentially making consumers worse off. For the equilibrium with overlapping products to occur, search costs must be sufficiently high. In the second type of search models, the multi-product search

[8]As discussed in the next section, there is also a large empirical literature investigating price discrimination in markets where search costs matter, with a focus on price patterns.

[9]See Ershov (2017) for an empirical application. Wildenbeest (2011) allows for vertical differentiated products but, unlike us, assumes that all consumers have the same preference for quality; hence, there is no scope for price discrimination. He finds that all firms use the same symmetric mixed strategy in utility space, meaning that firms use asymmetric price distributions depending on their product quality. In contrast, we find that firms might use different pricing strategies for the same product, with asymmetry arising because of price discrimination within the store.

[10]A recent strand of papers in the ordered search literature analyzes obfuscation by a multi-product firm monopolist (e.g., Gamp, 2019; Petrikaite, 2018). See Armstrong (2017) for a discussion.

[11]One-stop shopping considerations are also the driving force behind the evidence of price dispersion across stores documented by Kaplan *et al.* (2019).

assumption implies that discrimination is based on heterogeneity in consumers' shopping costs or some complementarity across the goods, which become the main determinants of firms' product choices (Klemperer, 1992).

Within this literature, Zhou (2014) finds that multi-product firms tend to charge lower prices than do single-product firms. This is driven not by the interaction between competition and price discrimination, as in our paper, but rather by a 'joint search' effect, i.e., multi-product firms charge less because they gain more by discouraging consumers from searching for competitors (see also McAfee, 1995). In Rhodes and Zhou (2019), increases in search costs imply that consumers value one-stop shopping more, thus making it more likely that the equilibrium involves multi-product firms. Unlike us, for small search costs, Rhodes and Zhou (2019) predict asymmetric market structures, with single-product and multi-product firms coexisting. The driving force underlying each of our predictions differs markedly: since in our model, consumers buy a single good, the multi-product firm equilibrium is driven not by one-stop shopping considerations, but rather by firms' incentives to price discriminate consumers with heterogeneous quality preferences.

Our model shares some aspects of Shelegia's (2012); notably, the fact that some consumers are informed about firms' prices, while others are not. However, unlike us, he assumes that consumers buy more than one good, and he does not analyze endogenous product choices. In the case of complements, Shelegia (2012) finds that prices are negatively correlated across goods in order to satisfy captive consumers' willingness to pay for the bundle. In our model, the positive price correlation across goods is driven instead by incentive compatibility considerations.

Like us, Garrett *et al.* (2019) introduce frictions in a model of price competition in which firms can carry more than one product but consumers buy only one.[12] The main difference between Garrett *et al.* (2019) analysis and ours is that they let firms decide qualities and prices *simultaneously*, while we model those choices as sequential. The simultaneous timing is appropriate in settings where firms can change product design rather quickly, or alternatively, when firms commit to prices for long periods of time, such as under long-term contracts. Sequential timing is better suited to capturing the notion that, in many markets, firms can change prices at will, while changes in product lines, which usually involve changes in production and/or retail facilities (Brander and Eaton,

[12]Another set of related papers analyzes pricing for add-ons. Ellison (2005) and Verboven (1999) consider models in which consumers are well informed about base product prices but don't know the price of add-ons unless they search. Critically, in these models the customers more likely to buy add-ons are also less likely to search. Our model is not one of add-on pricing, because some consumers are informed (i.e., they observe all prices) and others are not (i.e., they observe only prices of the store they visit). Furthermore, our model applies symmetrically for both products regardless of their quality.

1984), occur less often. This distinction is relevant, as in simultaneous settings firms cannot affect competition by pre-committing to quality choices, which is a fundamental driving force of our results.

Finally, our paper also relates to the literature that analyzes quality choices followed by imperfect competition, either quantity competition (Gal-Or, 1983; Wernerfelt, 1986; Johnson and Myatt 2003) or price competition with horizontal differentiation (Gilbert and Matutes, 1993; Stole, 1995; Kuksov, 2004). While one may view information frictions as equivalent to other forms of imperfect competition, they are not. In models of imperfect competition, for the equilibrium with overlapping (i.e., symmetric) quality choices to exist, competition must be sufficiently weak, e.g., as shown by Gal-Or (1983); under Cournot competition, the number of firms has to be sufficiently small. The same insight applies to models of price competition with horizontal product differentiation (e.g., Wernerfelt, 1986; Kuksov, 2004). In contrast, the impacts of information frictions on product line choices are different. Even if the number of uninformed consumers is arbitrarily small, firms do not have incentives to deviate from the equilibrium with overlapping product choices. The reason is that information frictions restore firms' monopoly power over uninformed consumers, even when competition for informed consumers (or *shoppers*) is fierce. This conclusion remains valid regardless of whether uninformed consumers visit a firm at random or whether visit the one that gives them higher ex-ante utility.[13]

The rest of this paper is organized as follows. Section 2 offers an example that conveys the main intuition of the model, while providing descriptive empirical evidence. Section 3 describes the model. Section 4 shows that in the absence of information frictions, firms can escape the Bertrand paradox by carrying non-overlapping product lines. In contrast, Section 5 shows that, if the costs of providing quality are sufficiently convex, arbitrarily small information frictions are enough to induce firms to choose overlapping product lines, even if this drives prices close to marginal costs. Section 6 characterizes equilibrium pricing for all potential product choice configurations, as well as the Subgame Perfect Equilibrium product choices for all levels of information frictions. Section 7 discusses the model's robustness to several extensions. Section 8 concludes. Selected proofs are provided in the appendix.[14]

[13]Indeed, we show that directed search by uninformed consumers strengthens our results, as firms have more reasons to become multi-product, compared to the case of non-observable product choices and uninformed consumers deciding randomly which firm to visit. See Section 7.

[14]Proofs of the characterization of pricing equilibria in subgames with asymmetric product choices are relegated to the online appendix. They follow a logic similar to proofs contained there in the appendix.

2 A Motivating Example

Price discrimination is pervasive in a wide range of markets in which information or search frictions matter. In gasoline markets, consumers have the choice of paying for full-service or self-service gasoline at the same station, or of searching for competing stations (Shepard, 1991). In the airline industry, travellers can choose to fly in business or economy class, or in economy class with certain restrictions (Borenstein and Rose, 1994). Other examples in which price discrimination, competition, and information frictions coexist include coffee shops (McManus, 2000), cereals (Nevo and Wolfram, 2002), theaters (Leslie, 2004), Yellow Pages advertising (Busse and Rysman, 2005), mobile telephony (Miravete and Röller, 2004), and cable TV (Crawford and Shum, 2007).

To build intuition on the main forces underlying our model, we focus on a market that fits well our modeling framework: online books.[15] While previous empirical papers have analyzed search in such markets (De Los Santos *et al.*, 2012; Hong and Shum, 2006), their focus has been on estimating buyers' search behavior for given product choices and prices. Our focus here is simply to motivate and illustrate predictions of the model by exploring firms' product choices and prices given consumers' search behavior. For this purpose, we analyze data collected daily for book prices at Amazon and Barnes & Noble, the two leading online booksellers, from December 2016 to March 2017, for each of the 2012-2016 #1 New York Times fiction and non-fiction best-sellers.[16]

2.1 Theoretical intuition

To develop intuition, let us think of two online stores competing to sell books to consumers with heterogeneous preferences for quality. Before choosing prices, booksellers must decide whether to offer both the hardcover and paperback editions of each book, or just one of the two, if any. Since the hardcover version is generally regarded of better quality than the paperback version, we will sometimes refer to the two as high- and low-quality goods, respectively. In turn, we will refer to those consumers willing to pay the extra cost of producing a hardcover as "high" types, and the remaining consumers as "low" types.

[15]Admittedly, our simple theoretical model does not capture all ingredients of real-world online markets. Notably, it assumes that shares of loyal consumers are symmetric across firms, which is not likely the case. Nevertheless, even with an asymmetric customer basis, the key properties of our equilibrium would be preserved: price dispersion, and at least one of the two firms would earn monopoly profits from its loyal consumers. See Narasimhan (1988) for an analysis of Varian's (1980) model with asymmetric shares of loyal consumers.

[16]We are grateful to Mar Reguant for providing these data.

In the absence of frictions, there are two types of equilibria. On the one hand, if the two stores offer both book editions, Bertrand competition would drive prices down to marginal costs. Because of the Bertrand reasoning, stores cannot deviate from this equilibrium by dropping one edition, as their profits would be zero in any event. On the other hand, CR's prediction is that firms can escape the Bertrand paradox by differentiating their product offerings. Indeed, it is also an equilibrium for one store to offer the hardcover edition, and the other store, the paperback. If one bookstore deviates from this equilibrium by carrying an additional format, competition would drive its price down to marginal costs, making such a deviation unprofitable. Furthermore, if the cost difference between hardcover and paperback is not too large, the store would have to give a discount on the other format to stop consumers from buying the one priced at marginal costs, further reducing the profitability of such a deviation.

Is CR's prediction robust to adding information frictions? To shed light on this question, suppose an arbitrarily small fraction of consumers visit one of two sites at random without searching further. These consumers are uninformed, as they observe only the version(s) of the book and price(s) of the site they visit.[17] If the site offers both book versions, uninformed consumers buy the one that gives them higher utility (if positive), given their quality preferences. If the site offers only one of the two versions, consumers buy that, as long as it gives them positive utility.

Following CR's prediction, suppose that the two stores offer different editions of the book. Now, the one carrying paperbacks might have incentives to also carry the hardcover edition. By doing so, it would profit more from the uninformed high types, as they willingly pay more for the hardcover. In turn, if costs of the hardcover relative to the paperback are sufficiently large, low types would not be willing to buy the hardcover even if that were sold at cost. Hence, since carrying the hardcover would not intensify competition for low types, profits made on the paperback would remain unchanged. It follows that, when the costs of providing high quality are sufficiently large, even an infinitesimally small amount of information frictions would be enough to rule out the equilibrium with non-overlapping products. This would also hold true under smaller quality differences, as long as the fraction of uninformed consumers is sufficiently large.

Alternatively, consider the Bertrand-like equilibrium at which both stores sell the hardcover and paperback editions. Now, bookstores face a trade-off when setting prices: on the one hand, they want to set high prices to maximize profits from selling books

[17] In the context of online books, De los Santos *et al.* (2012) show that, within a 7-day window, 76% of consumers visit only one store. They also report loyal consumers: 24% of consumers engage in multiple transactions but only buy from one store, even if that exhibits a higher price, suggesting specific store preferences, independent of prices.

to uninformed consumers; on the other hand, they want to set low prices to compete for the informed consumers who visit both sites. These countervailing incentives imply that the equilibrium must be in mixed strategies, with bookstores choosing random prices between the monopoly level and somewhere above marginal costs. Therefore, in equilibrium, bookstores' profits are the same *as if* they were monopolists over the uninformed consumers, but also *as if* competition washed away all profits from the informed consumers.

This has a key implication: bookstores' product choices are driven only by their incentives to better discriminate the uninformed consumers. Accordingly, they do not have incentives to drop either book format, as doing so would not enhance their market power over informed consumers, but rather would rather reduce the rents they can extract from uninformed consumers. While this incentive structure mimics a monopolist's, there is a fundamental difference with regard to duopolists': product overlap among competitors reduces their profits, to the extent that these go to zero as the fraction of uninformed consumers vanishes.

2.2 Evidence in the data

Product choice and pricing patterns observed in the online books' data are in stark contrast to CR's predictions, but can be rationalized after accounting for information frictions, as explained above. First, we find that stores sell both hardcover and paperback editions whenever both versions exist.[18] Second, we find that prices of hardcover and paperback editions of the same title do not remain constant. Consistent with our model, uninformed consumers make it profitable for stores to carry overlapping versions, which generates equilibrium price dispersion.[19] Third, we find that hardcovers are discounted more heavily than paperbacks during periods of sales.

Figure 1 provides evidence of price dispersion within and across stores. It plots a representative example of how prices (in US dollars) of a given title vary over time and differ across formats (hardcovers tend to be more expensive than paperbacks) as well as across stores. Beyond this anecdotal evidence, prices fluctuate substantially for each title even after removing book-store-format means. As shown in Figure 2a, this dispersion

[18]Of the 200 book titles that we consider, some are available only in paperback or in hardcover. All those available in both versions are carried by both Amazon and Barnes & Noble except for one instance, in which Amazon no longer offered the hardcover edition. Given the almost complete overlap, we focus our analysis on cases in which both stores carry both editions.

[19]Seim and Sinkinson (2016) provide evidence of mixed strategy pricing in online markets. Arguably, other reasons could also explain price dispersion in these markets, e.g., the use of algorithmic pricing (Chen *et al.*, 2016).

Figure 1: An example showing the price evolution of a given title

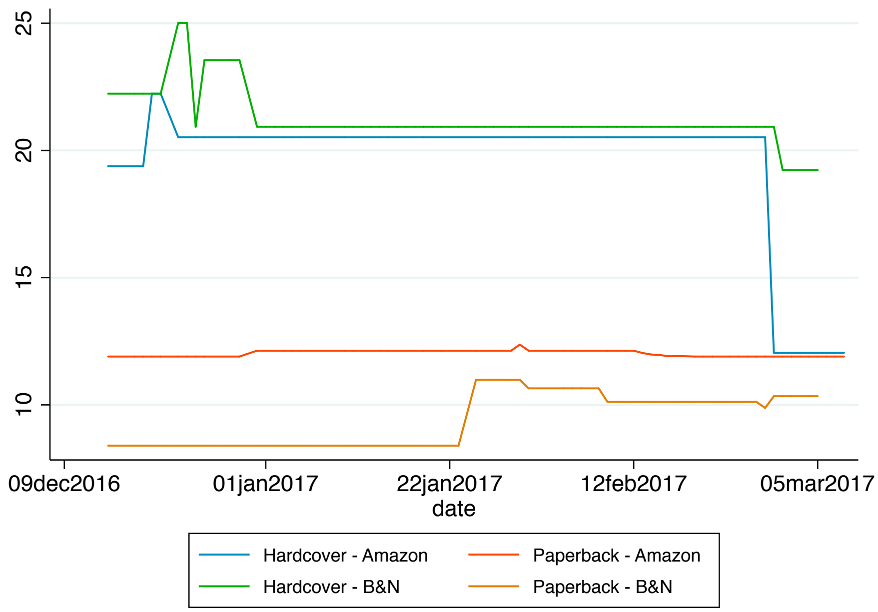

is not explained by common fluctuations, e.g., price fluctuations for particular books over time that are common across stores. Indeed, price differences across stores tend to fluctuate substantially, even after removing constant mean differences by book-format.

The interaction between price discrimination and information frictions can also explain the dispersion in relative prices of the two formats. Whereas existing search models cannot capture fluctuations in relative prices because those other models do not allow for price discrimination, our model predicts that information frictions lead not only to price differences across stores, but also to price differences across different versions of a book sold within a store, and over time. Figure 2b shows the distribution of relative prices between hardcover and paperback editions of each title. One can see substantial variation in relative prices, due partly to differences across titles and partly to variation in such relative prices over time. Furthermore, it shows that relative prices tend to be lower during periods of low prices (or *sales*). Conditioning on sales periods, average discounts on hardcovers and paperbacks (relative to the book-format-store means) are 10.2% and 6.9%, respectively. If we take into account within title and within store differences, hardcovers are more heavily discounted, 6.8 percentage points more than paperbacks. In levels, the price gap between the two formats shrinks by more than 2 USD during sales.[20]

In sum, the market for online books is characterized by a series of stylized facts. First, the norm is that all booksellers offer both hardcover and paperback versions of

[20]The online appendix contains more details of the empirical analysis.

Figure 2: Price patterns for online bookstores

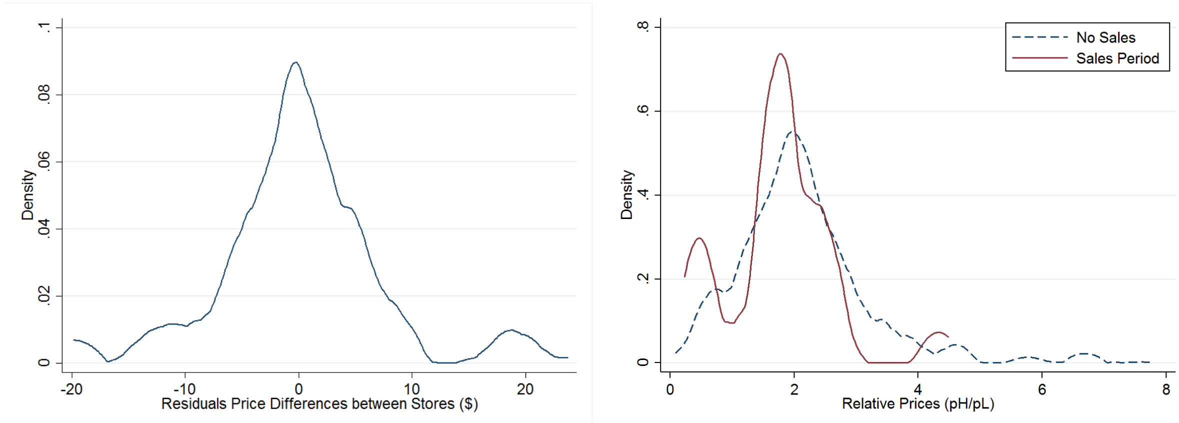

(a) Residual price differences between stores (b) Relative prices: sales vs. no sales

Notes: Panel (a) plots residuals from regressing the difference in prices at a given date between stores for the same book-format. Panel (b) shows hardcover prices relative to paperback prices during sales and no sales periods. Sales are defined as periods that exhibit discounts on both book editions that fall within the 25% (discount) quantile for a given store.

the same title, whenever available, even if this triggers intense competition for almost identical goods (up to the horizontal differences that consumers may perceive across stores). Additionally, book prices fluctuate substantially at the book-store level, but, more importantly, also across stores, making search meaningful. Relative prices between book versions also exhibit substantial dispersion, indicating another dimension firms use to sort out consumers and attract them from rivals. The model that we present next is capable of generating these predictions by highlighting the impact of information frictions on equilibrium product choices and price patterns.

3 The Model

3.1 Model Description

Consider a market served by two competing firms (sometimes referred to here as *stores*), which carry one or two goods: either a good with high quality q^H and high costs c^H, or another good with lower quality q^L and lower costs c^L, or both goods.[21] We use

[21]Without substantial effort, our model could be interpreted as one of quantity discounts, with firms offering the different quantities of the same product to consumers with either low or high demands. Results would go through as long as costs are not linear in the quality; for instance, if bigger bundles require costly product design features, such as packaging.

$\Delta q \equiv q^H - q^L > 0$ and $\Delta c \equiv c^H - c^L > 0$ to denote the quality and cost differences across goods.[22]

There is a unit mass of consumers who buy at most one good. Consumers differ in their quality preferences. A fraction $\lambda \in (0,1)$ of consumers have a low valuation for quality θ^L, while the remaining $(1-\lambda)$ fraction have a high valuation for quality θ^H, with $\Delta\theta \equiv \theta^H - \theta^L > 0$. As in Mussa and Rosen (1978), a consumer of type $i = L, H$ who purchases good $j = L, H$ at price p^j obtains net utility $u^i = \theta^i q^j - p^j$. We assume that the gross utility of a low (or high) type from consuming the low- (or high-) quality product always exceeds the costs of producing it, i.e., $c^i < \theta^i q^i$ for $i = L, H$. Therefore, for a consumer of type θ^i to be willing to buy good of quality q^i, the following incentive compatibility constraints must be satisfied

$$\theta^i q^i - p^i \geq \theta^i q^j - p^j, \tag{IC^i}$$

for $i, j \in \{L, H\}$ and $i \neq j$, which can also be re-written as

$$p^i \leq \theta^i q^i - \left(\theta^i q^j - p^j\right).$$

The second term on the right-hand side of the inequality represents consumers' *information rents*, i.e., the minimum surplus a consumer of type i needs to obtain to be willing to buy good i instead of good j.

The timing of the game is as follows. First, firms simultaneously decide which product(s) to carry (i.e., their "product line"). These choices are observed by firms but not by consumers. Second, firms simultaneously choose the prices of the product(s) they carry and consumers visit the stores to learn their product choices and respective prices. We will write (ϕ_i, ϕ_j) to denote firms' product choices, with $\phi_i \in \{\varnothing, L, H, LH\}$, and use $\Pi(\phi_i, \phi_j)$ to denote the profits of firm i at the pricing stage given those product choices.

Following Varian (1980), we assume that there is a fraction $\mu \leq 1$ of consumers who always visit the two stores and are therefore *informed* about where to find the cheapest product of each quality type. Since the remaining $1-\mu$ fraction of consumers visit only one store (with equal probability),[23] they are *uninformed* about the products and prices offered at the rival store. Hence, they can compare only the prices of goods sold *within*

[22] We can think of these costs as the wholesale prices at which retailers buy their products from either competitive manufacturers or a monopoly manufacturer. Endogenizing product qualities or the costs faced by retailers is beyond the scope of this paper.

[23] In some settings, it may be reasonable to assume that these consumers are uninformed about prices, but not about the product lines. Accordingly, we have also considered the case in which these consumers visit the store that gives them higher expected utility (and split randomly between the two stores in case of symmetry). Such considerations strengthen the main results of the paper. See Section 7.

the store they visit, not *across* stores. The fractions μ and λ are uncorrelated.[24] Once consumers have visited the store(s), they buy the product that gives them higher utility, provided that is non-negative. In case of indifference, low (high) type consumers buy the low (high) quality product. In what follows, we will use the fraction of uninformed consumers $1-\mu$ as a proxy for information frictions. Accordingly, the higher μ is, the lower the information frictions, with $\mu = 1$ representing a frictionless market.[25]

Assumptions In order to make the analysis meaningful, we rely on two assumptions that are standard in models of second-degree price discrimination (Tirole, 1988). The first guarantees that a monopolist carrying both goods finds it optimal to sort consumers. For the multi-product monopolist, the incentive compatible (i.e., constrained monopoly) prices are thus

$$\begin{aligned} p^L &= \theta^L q^L \text{ and} \\ p^H &= \theta^H q^H - \Delta\theta q^L \\ &= \theta^L q^L + \theta^H \Delta q. \end{aligned}$$

The alternative for the monopolist is to sell only good H to the high types at the (unconstrained) monopoly price $\theta^H q^H$, thus avoiding leaving information rents to the high types but also giving up profits on good L.[26] To guarantee that this alternative is indeed less profitable than selling the two goods requires that the profit from selling good L to low types be enough to compensate for the information rents that must be left with the high types:

$(A1)$ $\lambda(\theta^L q^L - c^L) \geq (1-\lambda)\Delta\theta q^L.$

Note that $(A1)$ is evaluated at monopoly prices. Assuming that a monopolist prefers to carry all qualities does not necessarily imply that the same holds true when competition drives prices below the monopoly level.

Our second assumption guarantees no 'bunching' at the competitive solution. This requires marginal cost pricing to be incentive compatible, which is equivalent to assuming that the high types are willing to pay for the extra cost of high quality, whereas the low types are not:

[24] As discussed in Section 7, our main results do not change if we allow μ and λ to be correlated.

[25] In the online appendix, we endogenize μ following the fixed-sample search model of Burdett and Judd (1983), but allow consumers to differ in terms of their search costs.

[26] Note that this alternative assumes that serving the high types with product H is more profitable than serving all consumers with product H at price $\theta^L q^H$. This is guaranteed by our assumption $(A3)$ below.

(*A*2) $\Delta c \in \left(\theta^L \Delta q, \theta^H \Delta q\right)$.

Implicit in (*A*2) is the standard property that the cost of providing quality must be strictly convex in quality, i.e., $c^H/q^H > c^L/q^L$; otherwise, either type would buy the high-quality product or nothing at all (CR adopt a similar assumption).[27]

Finally, in order to reduce the number of cases we need to consider without affecting results, we will assume that a monopolist carrying only good H prefers to extract all the surplus from the high types, even if that implies not selling to the low types, which would require reducing the price to $\theta^L q^H$:[28]

(*A*3) $(1-\lambda)(\theta^H q^H - c^H) \geq \theta^L q^H - c^H$.

It follows that the single-product monopoly prices are $\theta^H q^H$ for the firm carrying good H and, as implied by (*A*1), $\theta^L q^L$ for the firm carrying good L. In what follows we will denote the single-product monopoly profits as $\pi^i \equiv \theta^i q^i - c^i$ for $i \in \{L, H\}$.

Minmax profits Inspection of assumption (*A*2) above allows us to obtain useful expressions for the analysis that follows. As implied in (*A*2), the maximum profits that can be made from product $i \in \{L, H\}$ when good $j \neq i$ is priced at marginal costs are strictly positive. Since firms would never sell their products below marginal costs, these constitute *minmax profits*. In particular, if good L is sold at c^L, good H can at most be sold at the highest price that satisfies the high types' incentive compatibility constraint, i.e., $p^H \leq c^L + \theta^H \Delta q$. This gives per unit profits of

$$\varphi^H \equiv \theta^H \Delta q - \Delta c > 0.$$

The minmax profits for good H are always strictly below monopoly profits π^H given that, for all values of c^L, good L imposes a competitive constraint on good H.

In turn, if good H is sold at c^H, good L can at most be sold at the highest price that satisfies the low types' participation and incentive compatibility constraints, i.e., $p^L \leq \min\left\{\theta^L q^L, c^H - \theta^L \Delta q\right\}$. This gives per unit profits of

$$\varphi^L \equiv \min\left\{\pi^L, \Delta c - \theta^L \Delta q\right\} > 0.$$

For $c^H \geq \theta^L q^H$, the participation constraint binds first, so good L can be sold at the monopoly price even when good H is priced at marginal cost, i.e., $\varphi^L = \pi^L$. Alternatively,

[27]Note also that convexity ensures a non-empty region of λ values for which (*A*1) and (*A*2) are valid.

[28]Note that this assumption is redundant when the costs of providing high quality are large enough, i.e., $c^H \geq \theta^L q^H$, but it does imply that these costs cannot be much lower than $\theta^L q^H$.

for $c^H < \theta^L q^H$, the incentive compatibility constraint binds first, so the minmax profits for good L are strictly below monopoly profits, i.e., $\varphi^L = \Delta c - \theta^L \Delta q < \pi^L$.

In sum, the per unit profits that a firm monopolizing good $i \in \{L, H\}$ loses when product $j \neq i$ is made available at marginal cost equal $\pi^i - \varphi^i \geq 0$ (with equality only for good L when $c^H \geq \theta^L q^H$). We will sometimes express profits as functions of φ^H and φ^L. The following equalities will be particularly useful throughout the analysis: $\pi^H - \varphi^H = \pi^L + \Delta\theta q^L$, and if $c^H < \theta^L q^H$, then $\pi^L - \varphi^L = \theta^L q^H - c^H$.

We are now ready to solve the game. We start by analyzing the case in which all consumers are informed, $\mu = 1$, then introduce an arbitrarily small fraction of uninformed consumers, $\mu \to 1$, and finish by providing a full equilibrium characterization for all $\mu \in [0, 1)$.

4 Escaping the Bertrand Paradox

In this section we characterize the equilibrium of the game under no information frictions.

Proposition 1 *Assume all consumers are informed, $\mu = 1$. There exist two (pure strategy) Subgame Perfect Equilibria (SPE):*

(i) The "overlapping" equilibrium (LH, LH), at which both firms make zero profits.

(ii) The "specialization" equilibrium (L, H), at which both firms make strictly positive profits.[29]

Proof. See the appendix. ■

In the absence of information frictions, two types of equilibria exist: (i) a Bertrand equilibrium in which firms carry both products and make zero profits ("overlapping equilibrium"), and (ii) an equilibrium in which firms carry non-overlapping product lines, each of which makes a strictly positive profit ("specialization equilibrium"). Hence, in the absence of search costs, simultaneous quality choices followed by price competition allow firms to escape the Bertrand paradox.[30]

To understand why the latter equilibrium exists, first note that under product choices (L, H), a pure strategy equilibrium does not exist at the pricing stage. This stems from an important result: in equilibrium, firms' prices must satisfy incentive compatibility.

[29] If $c^H < \theta^L q^H$, a symmetric mixed strategy equilibrium also exists with positive profits such that firms choose L and H with positive probability. Otherwise, if $c^H \geq \theta^L q^H$, this equilibrium does not exist, as it is dominated by playing LH.

[30] Firms would also escape the Bertrand paradox if one of them carries both products and the other carries none. This equilibrium is not only uninteresting but also irrelevant in our analysis, as it disappears as soon as we introduce information frictions.

Otherwise, the firm selling good H would sell nothing and would thus be better off reducing its price to satisfy incentive compatibility. However, if the high types' incentive compatibility constraint is binding, the firm carrying good L could in turn attract all customers by slightly reducing its own price. Since these opposing forces destroy any candidate price choice in pure strategies, the equilibrium has to be in mixed strategies. Furthermore, all prices in the support of the mixed strategies must be strictly above marginal costs.

This has meaningful implications for equilibrium product choices. First, since at (L, H) product L is priced *above* marginal costs, profits on good H are strictly above its minmax. If firm H deviated to also carrying good L, p^L would be driven down to marginal costs. Hence, profits on good L would be zero and the profits on good H would be driven down to its minmax, making such a deviation unprofitable. Similarly, since at (L, H) product H is priced *above* marginal costs, profits on good L are (weakly) above its minmax. If firm L deviated to also carrying good H, it would make no profits on good H and would (weakly) reduce its profits on good L as competition for good H becomes fiercer.[31] Finally, if either firm deviated so that the two products overlapped, leading to (L, L) or (H, H), they would both make zero profits. In sum, since firms do not gain by deviating from (L, H), the "specialization equilibrium" constitutes a SPE of the game under no information frictions.

CR disregard the "overlapping equilibrium" by requiring that both firms make strictly positive profits in equilibrium. One way to justify this choice would be to assume that firms face (even infinitesimally small) fixed costs of carrying a product. Another way would be to rely on the Pareto criterion, as both firms make strictly higher profits at the "specialization equilibrium".[32] CR's focus on the "specialization equilibrium" has been very influential in disseminating the view that firms can soften competition by differentiating their product choices. The next section, however, shows that CR's prediction is not robust to introducing information frictions.

[31] If $c^H \geq \theta^L q^H$, the firm carrying good L makes the same profits at (L, H) as at (LH, H) since good H imposes no competitive constraint on good H. In any event, firm L could increase its profits by carrying good H, as well.

[32] However, as will be seen in the next section, one compelling reason to focus on the "overlapping equilibrium" is that it naturally converges to the Bertrand equilibrium as information frictions vanish, while the "specialization equilibrium" may not.

5 Back to the Bertrand Paradox

Before solving the game for all $\mu \in [0, 1)$, in this section we show that arbitrarily small information frictions, $\mu \to 1$, are enough to give rise to an equilibrium with overlapping product lines and positive profits. Furthermore, we show that if the costs of providing high quality are large enough, the "specialization" equilibrium ceases to exist.

To explore this phenomenon in more detail, let us first analyze pricing incentives at the subgame with "overlapping" product choices (LH, LH). Information frictions, no matter how small, imply that marginal cost pricing is not an equilibrium, as firms could make positive profits from uninformed consumers. Similarly, setting prices at the (constrained) monopoly level is not an equilibrium either, as firms would have incentives to charge slightly lower prices to attract informed consumers. More generally, information frictions rule out any equilibrium candidate in pure strategies, as firms face a trade-off between charging high prices to exploit uninformed consumers versus charging low prices to attract informed consumers. Since firms must be indifferent to charging any price vector in the support, expected equilibrium profits can be computed by summing the profits of each good at the upper bound, where firms optimally serve their share of uninformed consumers at (constrained) monopoly prices,

$$\Pi(LH, LH) = \frac{1-\mu}{2} \left[\lambda \pi^L + (1-\lambda)(\pi^H - \Delta\theta q^L)\right]. \tag{1}$$

Importantly, each firm's equilibrium profits are a fraction $(1-\mu)/2$ of the multi-product monopolist's profits because, at the upper bound, firms profit only from uninformed consumers. For prices below the upper bound, firms make the same profits in expectation: the positive profits they obtain from informed consumers compensate for lower profits obtained from uninformed consumers. As μ approaches 1 and all customers become informed, the equilibrium price distributions concentrate around marginal costs, and firms' profits are driven down to nearly zero. The Bertrand outcome is thus restored.

Could firms escape the Bertrand paradox by having one drop a product, either L or H?[33] Let us first analyze incentives for moving from (LH, LH) to (H, LH). Since a pure strategy equilibrium does not exist, and firms have to be indifferent across all prices in the support, expected profits for product H equal those from serving uninformed consumers at the upper bound. Since firm H is not constrained by incentive compatibility, according to $(A3)$, its optimal price at the upper bound is the (unconstrained) monopoly price. Its expected profits become

$$\Pi(H, LH) = \frac{1-\mu}{2}(1-\lambda)\pi^H. \tag{2}$$

[33]Neither firm has incentives to drop both products altogether, as they both make positive profits at (LH, LH).

Since firm H's profits are a fraction $(1-\mu)/2$ of monopoly profits, comparing (1) and (2) is equivalent to assessing the monopolist's incentives for carrying the high-quality good only versus carrying both goods. Assumption $(A1)$ guarantees that (1) exceeds (2), as the losses from not selling the low-quality product exceed the information rents left to the high types. Thus, even though product L erodes the rents made on product H, the firm is better off carrying it.

The alternative is for one of the two firms to drop product H, thus moving from (LH, LH) to (L, LH). Now the expected profits of firm L must equal the profits from serving all uninformed consumers at the unconstrained monopoly price,[34]

$$\Pi(L, LH) = \frac{1-\mu}{2}\pi^L,$$

which again is a fraction $(1-\mu)/2$ of monopoly profits. By $(A2)$, this payoff is strictly less than (1) since the firm gives up the extra profit that firm L could make by selling the high-quality good to the uninformed high types, who are willing to pay for the extra cost of providing quality.

In sum, firms' profits are the same *as if* they exploited their monopoly power over uninformed consumers and competed fiercely for the informed ones, obtaining no profits from the latter. Hence, firms' incentives to price discriminate through product choice mimic those of the monopolist. Consequently, in the presence of arbitrarily small information frictions, a SPE exists with overlapping product lines (LH, LH) in which firms make strictly positive profits, in contrast to CR's prediction.

To assess whether this equilibrium is unique or not, let us first note that the "specialization" equilibrium of Proposition 1 is ruled out when the cost of providing high quality is sufficiently large, $c^H \geq \theta^L q^H$ (or equivalently, when the costs of providing quality is sufficiently convex). Starting at (L, H), firm L is strictly better off adding product H given that under (LH, H) it can now price discriminate uninformed consumers without eroding its profits on good L. Indeed, the firm would be able to increase its profits by $(1-\mu)(1-\lambda)\varphi^H/2 > 0$ from selling the high- rather than the low-quality product to uninformed high types, while it would still make profits $\lambda\pi^L$ from low types. Intuitively, low types would never want to buy the high-quality good even if it were sold at cost.

In contrast, if the costs of high quality are sufficiently low, the addition of good H erodes the rents of good L, making firm L worse off: the rents on product H are infinitesimally small while the profits on good L would go down by $\lambda\left(\pi^L - \varphi^L\right) > 0$. Similarly, firm H does not want to add product L as its profits would fall by $(1-\lambda)\left(\pi^H - \varphi^H\right) > 0$.

[34] Note that in this case, the firm would serve both low and high types, since the latter are also willing to buy the low-quality product at the (unconstrained) monopoly price for product L.

Thus, the "specialization" equilibrium survives the introduction of infinitesimally small information frictions *only* when the costs of providing high quality are sufficiently low.

Our second Proposition summarizes these results.

Proposition 2 *Assume that the mass of uninformed consumers is infinitesimally small,* $\mu \to 1$.

(i) The "overlapping" equilibrium (LH, LH) *constitutes a SPE for all parameter values. Equilibrium prices approximate marginal costs.*

(ii) The "specialization" equilibrium (L, H) *constitutes a SPE if and only if* $c^H < \theta^L q^H$. *Equilibrium prices are strictly above marginal costs.*

Proof. See the discussion above. A formal derivation can be found as a particular case of the proof to Proposition 7. ■

The addition of even infinitesimally small information frictions implies that, when the costs of providing high quality are high enough (or when low types do not value high quality sufficiently), firms can no longer escape the Bertrand paradox by differentiating their product lines. Indeed, the "specialization" equilibrium no longer exists, making the Bertrand-like "overlapping" equilibrium the unique SPE of the game.

In our model, the "overlapping" equilibrium always exists. This is in contrast to previous papers analyzing quality choices followed by imperfect competition (Gal-Or, 1983; Gilbert and Matutes, 1993; Johnson and Myatt, 2003; Stole, 1995; Wernerfelt, 1986), in which the "overlapping" equilibrium exists *only if* the rents created by imperfect competition are high enough (e.g., few firms competing *à la* Cournot). In those papers, just as in CR, there is a tension between competition and price discrimination: competition reduces the rents on overlapping products while enlarging consumers' information rents, thus reducing the gains from price discrimination.

In this paper, under the "overlapping" equilibrium that arises with information frictions, such tension is not present because firms only care about the profits made form uninformed consumers, from whom they obtain monopoly profits (in expectation). Thus, firms' product choices are driven solely by their incentives to discriminate consumers, leading them to carry the full product range even when rents created by information frictions are arbitrarily small. This shows that the impact of information frictions on product choices, and through these on prices, may differ from other forms of imperfect competition.

6 Equilibrium Product Lines and Prices

In this section, we characterize equilibrium product and price choices for all values of $\mu < 1$. We show that the "overlapping" equilibrium is robust to introducing information frictions, no matter how big or small. In contrast, the "specialization" equilibrium fails to exist when the mass of informed consumers μ is sufficiently small or, for all μ, when the cost of providing high quality c^H is sufficiently high. In general, the "overlapping" equilibrium is more likely to be unique for a smaller mass of informed consumers and/or higher costs of providing high quality.

We again proceed by backward induction, analyzing first equilibrium pricing behavior and then product choices. The pricing subgames also help us understand pricing decisions for non-overlapping product configurations, which may prove relevant to cases in which product choices are constrained by factors outside our model (e.g., fixed costs of carrying a product).

6.1 Pricing Behavior

We first provide an important property of multi-product firms' pricing behavior.

Lemma 1 *In equilibrium, multi-product firms choose incentive-compatible prices for their products, i.e.,* $\Delta p \in \left[\theta^L \Delta q, \theta^H \Delta q\right]$.

Proof. See the appendix. ■

The lemma above shows that choosing prices to satisfy incentive compatibility is always optimal for a multi-product firm. The intuition is simple. If the price of the high-quality product is so high that nobody buys it, reducing p^H while leaving p^L unchanged is profitable for the firm to attract high types and capture a larger profit margin on them without affecting low types' decisions. Similarly, if low-quality product's price is so high that nobody buys it, decreasing p^L while leaving p^H unchanged is profitable for the firm to attract low types and capture a larger margin on them without affecting high types' decisions. This result constitutes an important departure from Varian (1980), as it implies that the price of one product cannot be selected independently from the price of another product *within* the same store.[35]

We are now ready to characterize equilibrium pricing for every possible subgame.

[35] This result is in contrast to Johnson and Myatt's (2015) prediction. In a model of quality choice followed by Cournot competition, they find conditions under which the equilibrium prices chosen by multi-product oligopolists are close to the single-product prices.

Full product overlap We start by considering subgames with full product overlap: (LH, LH), (L, L), and (H, H). The next proposition characterizes equilibrium pricing under the former one.

Proposition 3 *Given product choices (LH, LH), the equilibrium must be in mixed strategies. In addition, any symmetric mixed-strategy equilibrium must satisfy the following properties:*

(i) Prices at the upper bound of the price support correspond to the (constrained) monopoly prices, $\overline{p}^L = \theta^L q^L$ and $\overline{p}^H = \theta^H q^H - \Delta\theta q^L = \theta^L q^L + \theta^H \Delta q$, so that the high types' incentive compatibility constraint is binding, $\Delta\overline{p} \equiv \overline{p}^H - \overline{p}^L = \theta^H \Delta q$.

(ii) Prices at the lower bound of the price support are strictly above marginal costs, $\underline{p}^i > c^i$ for $i = L, H$, and such that the high types' incentive compatibility constraint is not binding, $\Delta\underline{p} \equiv \underline{p}^H - \underline{p}^L < \theta^H \Delta q$.

(iii) Any pair of prices, $p^H \in [\underline{p}^H, \bar{p}^H]$ and $p^L \in [\underline{p}^L, \bar{p}^L]$, is chosen according to some joint distribution function $F^{LH}(p^H, p^L)$ that is consistent with Lemma 1: $p^H - p^L \equiv \Delta p \in [\theta^L \Delta q, \theta^H \Delta q]$.

Proof. See the appendix. ■

The non-existence of pure strategy equilibria is shared with most search models, starting with Varian (1980) (see also Burdett and Judd (1993) and McAfee (1995), among others). This situation stems from firms' countervailing incentives, as on the one hand, firms want to reduce prices to attract informed consumers, but on the other, they want to extract all rents from uninformed consumers.

Despite this similarity, our analysis shows that equilibrium pricing by multi-product firms has a distinctive feature: it is constrained by incentive compatibility (Lemma 1). This comes up clearly when characterizing the upper bound of the price support: firms are not able to extract all surplus from uninformed high types because they have to give up information rents $\Delta\theta q^L$.[36] Hence, because of incentive compatibility, firms profit less on the high-quality good than in the single-product case, in contrast to McAfee (1995).

Since firms make strictly positive profits at the upper bound, prices at the lower bound must be strictly above marginal costs. The reduction in prices from the upper to the lower bound is more pronounced for the high-quality product than for the low-quality one. Competition for high types is fiercer because selling the high-quality product is more profitable. In turn, this implies that at the lower bound, the incentive compatibility constraint for high types is not binding, so the price wedge between the two products at

[36] Note that, in equilibrium, nothing prevents $F^{LH}(p^H, p^L)$ from being such that a firm plays $\bar{p}^L$ together with $p^H \in [\underline{p}^L + \theta^L \Delta q, \bar{p}^H]$, or alternatively, $\underline{p}^L$ together with $p^H \in [\underline{p}^H, \underline{p}^L + \theta^H \Delta q]$.

the upper bound is wider than at the lower bound. We can conclude that high-quality products are relatively cheaper during "sales" periods *à la* Varian, i.e., when both goods are priced at the lower bounds of the price support. Even when firms do not price the two goods simultaneously at the lower bound, the relative price difference never exceeds that under monopoly, $\theta^H \Delta q$, as otherwise incentive compatibility would not be satisfied (Lemma 1). Thus, competition among multi-product firms reduces the relative prices of the two goods.

Since firms must be indifferent to charging any price in the support, including the upper bound, expected equilibrium profits are unambiguously given by

$$\Pi\left(LH, LH\right) = \frac{1-\mu}{2}\left[\lambda\pi^L + (1-\lambda)(\pi^H - \Delta\theta q^L)\right]. \tag{3}$$

As noted in the previous section, these profits are a fraction $(1-\mu)/2$ of the (constrained) monopoly profits.

At the lower bound, each firm attracts all informed consumers plus its share of uninformed consumers of each type. Hence, expected profits can also be expressed as a function of the lower bounds,

$$\Pi\left(LH, LH\right) = \frac{1+\mu}{2}\left[\lambda(\underline{p}^L - c^L) + (1-\lambda)(\underline{p}^H - c^H)\right]. \tag{4}$$

Since there are two goods and only one profit level, as given by equations (3) and (4), the problem has an extra degree of freedom: potentially many price pairs $\underline{p}^L > c^L$ and $\underline{p}^H > c^H$ satisfying $\Delta\underline{p} < \theta^H \Delta q$ yield the same equilibrium profits. This implies that, even though equilibrium profits are unique and well defined, there might be a multiplicity of mixed strategy equilibria.[37]

Finally, there could also be full overlap among single-product firms, (L, L) and (H, H). Since single-product firms selling the same product are not constrained by incentive compatibility, they play a mixed strategy equilibrium with an upper bound equal to the (unconstrained) monopoly price, as in Varian (1980). However, the presence of heterogeneous consumers would add a small twist to Varian's pricing. In particular, there could now be a gap in the price support between the prices at which firms are indifferent to serving high types only (at a high price) versus serving both types (at a lower price). Note that, in this case, low types are left out of the market with some positive probability. Other than this, since equilibrium profits are fully determined by the upper bound, equilibrium profits are as in Varian (1980).

[37] See McAfee (1995) and Shelegia (2012) for a similar result. Appendix B characterizes one equilibrium of this subgame.

Partial product overlap Let us now characterize equilibrium pricing in the subgames with partial overlap: (L, LH), (H, LH). Interestingly, even though the single-product firm does not face an incentive compatibility constraint *within* its store, its pricing is nevertheless affected by incentive compatibility considerations through the effect of competition *across* stores.

The following Proposition characterizes the price equilibrium at the (L, LH) subgame.

Proposition 4 *Given product choices (L, LH):*

(i) A pure strategy equilibrium does not exist.

(ii) At the unique mixed-strategy equilibrium, firm LH charges $p^H = p^L + \theta^H \Delta q$, and both firms choose p^L in $\left[\underline{p}^L, \theta^L q^L\right]$, with firm L putting a probability mass at the upper bound.

Proof. See the online appendix. ■

In equilibrium, the two firms choose random prices for the low-quality product over a common support. In turn, given its price choice for the low-quality good, the multi-product firm prices the high quality product to just barely comply with incentive compatibility for high types. Hence, unlike the previous case, the price difference between the two products remains constant at $\theta^H \Delta q$ over the whole support, and the density of prices for the high-quality product is the same as that for the low-quality product (just shifted out to the right by $\theta^H \Delta q$). It follows that, whenever the multi-product firm has the low price for the low-quality product, all informed consumers (of either low or high type) buy from it. Otherwise, if the single-product firm charges the lower price for the low-quality product, it serves all informed consumers, including both low and high types.[38] Its profits nevertheless are determined by its upper-bound price. As before, its profits are a fraction $(1-\mu)/2$ of the firm's monopoly profits

$$\Pi(L, LH) = \frac{1-\mu}{2} \pi^L.$$

Now we turn to characterizing the price equilibrium at the (H, LH) subgame.

Proposition 5 *Given product choices (H, LH), there exists $\hat{\mu} \in (0,1)$ such that:*

[38] It is worth noting that the multi-product firm charges lower prices, on average, compared to the single-product firm. The reason is that, when it has the lower price, its ability to discriminate between low and high types allows the multi-product firm to make extra profits $\mu(1-\lambda)\varphi^H$ from informed high types. Since the multi-product firm has stronger incentives to undercut its rival's price, the single-product firm must put a probability mass at the upper bound. In turn, since the two firms cannot put a mass at the same price, it follows that when the single-product firm is pricing at the upper bound, it is only selling to uninformed consumers with probability one.

(i) For $\mu \leq \hat{\mu}$, there exists a unique pure strategy equilibrium: firm H chooses the (unconstrained) monopoly price $p^H = \theta^H q^H$, and firm LH chooses the (constrained) monopoly prices, $p^H = \theta^H q^H - \Delta\theta q^L$ and $p^L = \theta^L q^L$.

(ii) For $\mu > \hat{\mu}$, a pure strategy equilibrium does not exist. In the mixed-strategy equilibrium, firm LH chooses prices p^H in $\left[\underline{p}^H, \theta^H q^H - \Delta\theta q^L\right]$ with a mass on its upper bound, and $p^L = \min\left\{\theta^L q^L, p^H - \theta^L \Delta q\right\}$. Firm H chooses prices p^H in $\left\{\left[\underline{p}^H, \theta^H q^H - \Delta\theta q^L\right], \theta^H q^H\right\}$ with a (strictly) positive mass on its upper bound.

Proof. See the online appendix. ■

There now exists a pure strategy equilibrium as long as the fraction of informed consumers μ is sufficiently small. At this equilibrium, the multi-product firm charges the (constrained) monopoly prices while the single-product firm charges the (unconstrained) monopoly price for the high-quality product.

When the fraction of informed consumers is higher, the above is no longer an equilibrium, as the single-product firm now has more reason to fight for informed consumers. In this case, the equilibrium must be in mixed strategies.[39,40] The precise shape of the mixed strategy equilibrium depends on whether or not it pays for firm H to serve low types.

If $c^H \geq \theta^L q^H$, it never pays firm H to serve low types because the costs of high quality exceed their willingness to pay for it. Thus, the two firms compete for informed high types only, while the low-quality product is still priced at the monopoly level, $\theta^L q^L$. Since the incentive compatibility constraint of the multi-product firm is not binding, its profits are the same as if the two products were sold independently. In contrast, when $c^H < \theta^L q^H$, low types might be tempted to buy the high-quality good when its price is sufficiently low. In this case, the price of the low-quality good must be reduced below its monopoly level to achieve separation.

Regarding the single-product firm, since $\theta^H q^H - \Delta\theta q^L$ is the highest price that the multi-product firm would ever charge for the high-quality good, the firm will play either the (unconstrained) monopoly price, $\theta^H q^H$, or something less than the (constrained) monopoly price, $\theta^H q^H - \Delta\theta q^L$. Any price in between is unprofitable, either because it does not extract enough from uninformed high types or because it does not attract informed consumers when the multi-product firm happens to price the good at or below

[39] Interestingly, there is continuity between the pure- and mixed-strategy equilibria. The two firms charge the upper bounds of their price supports, $\theta^H q^H - \Delta\theta q^L$ and $\theta^H q^H$, with positive and identical mass, which fades as μ grows larger—from one, when $\mu \to \hat{\mu}$ toward zero, when $\mu \to 1$.

[40] Unlike in subgame (LH, LH), the equilibrium is now unique: since one firm has only one product, there are no longer two degrees of freedom, as in the symmetric two-product case.

$\theta^H q^H - \Delta\theta q^L$. In either case, profits remain as in the pure strategy equilibrium because $\theta^H q^H$ always belongs to the price support. Therefore, for all μ,

$$\Pi(H, LH) = \frac{1-\mu}{2}(1-\lambda)\pi^H,$$

again a fraction $(1-\mu)/2$ of the monopoly profits.[41]

Non-overlap Let us now move to characterizing equilibrium pricing in the subgames with no product overlap: $(\varnothing, L)$, $(\varnothing, H)$, $(\varnothing, LH)$ and (L, H). The first three product configurations correspond to the single-product monopoly solution. Hence, we turn our attention to the more interesting subgame with specialized firms, (L, H).

Proposition 6 *Given product choices (L, H), there exists $\tilde{\mu} \in (\hat{\mu}, 1)$ such that:*

(i) For $\mu \leq \hat{\mu}$, there exists a unique pure strategy price equilibrium: firms charge the (unconstrained) monopoly prices $p^H = \theta^H q^H$ and $p^L = \theta^L q^L$.

(ii) For $\mu > \hat{\mu}$, a pure strategy equilibrium does not exist. At the unique mixed-strategy equilibrium, firm L chooses prices p^L in $\left[\underline{p}^L, \theta^L q^L\right]$ with a mass on the upper bound. If $\mu \in (\hat{\mu}, \tilde{\mu})$, firm H chooses prices p^H in $\left\{\left[\underline{p}^H, \theta^H q^H - \Delta\theta q^L\right], \theta^H q^H\right\}$ with a mass on the upper bound that falls to zero as $\mu \to \tilde{\mu}$; if $\mu \geq \tilde{\mu}$, $\theta^H q^H$ is not part of firm H's support.

Proof. See the online appendix. ■

Equilibrium pricing at subgames (L, H) and (H, LH) share some similarities. In particular, just as in Proposition 5, if the mass of informed consumers μ is small enough, there exists a pure strategy equilibrium as the firm selling the high-quality product is better off serving uninformed high types at the (unconstrained) monopoly price than competing for informed high types.[42] Furthermore, there is continuity between the pure- and mixed-strategy equilibria in that the probability mass that the high-quality firm puts on the (unconstrained) monopoly price fades away as μ grows larger.

The main difference between the two subgames is that, under (L, H), the high quality firm chooses not to include the unconstrained monopoly price in the support when μ is very large. The reason is that profits from serving a small fraction of uninformed consumers fall short of profits from fighting for informed consumers.[43]

[41] Just as in the previous subgame, the equilibrium price distribution used by the multi-product firm for the high-quality good (weakly) first-order stochastically dominates that of the single-product firm. It follows that, on average, the price charged by the single-product firm for the high-quality product exceeds the price charged by the multi-product firm.

[42] Note that the threshold for the existence of a pure-strategy equilibrium is the same under both subgames.

[43] At subgame (H, LH), competition for good H is more intense given that both firms carry it. This explains why in that case firm H always puts mass at the unconstrained monopoly price, while at subgame (H, L) firm H eventually decides not to include it in its price support.

6.2 Product Line Choices

We are now ready to analyze product line decisions given the continuation equilibria characterized above. For this purpose, it is useful to implicitly define the threshold μ^* as

$$(1-\mu^*)(1-\lambda)\left(\pi^H-\pi^L\right)=(1+\mu^*)\lambda\left(\pi^L-\varphi^L\right). \tag{5}$$

Note that μ^* is increasing in c^H, and that $\mu^*=1$ for $c^H \geq \theta^L q^H$. The following Proposition characterizes the Subgame Perfect Equilibrium (SPE) product choices.

Proposition 7 *(i) If $\mu<\mu^*$, the "overlapping" equilibrium (LH, LH) constitutes the unique SPE of the game. (ii) Otherwise, both the "overlapping" equilibrium (LH, LH) and the "specialization" equilibrium (L, H) constitute a SPE of the game.*

Proof. See the appendix. ■

In Proposition 2, we showed that a SPE with overlapping product lines exists in the presence of an arbitrarily small number of uninformed consumers. Proposition 7 now shows that this prediction remains valid for all values of μ. The underlying logic remains the same: the existence of the "overlapping" equilibrium hinges upon the incentives of firms to mimic those of a monopolist, regardless of whether μ is large or small.

Regarding the existence of the "specialization" equilibrium, Proposition 2 showed that it exists for $\mu \to 1$ as long as $c^H < \theta^L q^H$. Thus, the equilibrium with overlapping qualities is unique when information frictions are arbitrarily small, as long as the costs of providing high quality are large enough. Proposition 7 now shows that the presence of uninformed consumers relaxes the condition for the uniqueness of the equilibrium with overlapping qualities. In particular, whereas $c^H < \theta^L q^H$ is still needed to guarantee the existence of the specialization equilibrium, it no longer is sufficient: additionally, information frictions must be low enough for the gains from softening competition to exceed the costs of giving up profitable opportunities to discriminate. To see this in more detail, consider the incentives to deviate from the "specialization" equilibrium by the firm carrying product L. Adding product H would allow the firm to better discriminate high types, thus making extra profits $\left(\pi^H-\pi^L\right)$ from selling product H to uninformed high types with probability $(1-\mu)(1-\lambda)/2$. In contrast, adding product H would also intensify competition for product L, forcing the firm to give up rents $\left(\pi^L-\varphi^L\right)$ on all the low types (excluding the uninformed low types that visit the rival's store) with probability $(1+\mu)\lambda/2$. The magnitudes of the two effects coincide at $\mu=\mu^*$, as implicitly defined in equation (5). In turn, since in expectation firms benefit only from discriminating uninformed consumers, the softening of competition effect dominates the incentives to discriminate only when the mass of uninformed consumers $(1-\mu)$ is sufficiently small, i.e., when $\mu \geq \mu^*$. Therefore, for $\mu < \mu^*$, the "specialization" equilibrium breaks down, making the "overlapping" equilibrium the unique SPE of the game.

The fact that μ^* is increasing in c^H means that, as the cost of high-quality provision increases to $\theta^L q^H$, the set of μ values for which the "overlapping" equilibrium is unique is enlarged; beyond that level, the "overlapping" equilibrium is the unique SPE for all μ. If $c^H \geq \theta^L q^H$, $\mu^* = 1$, implying that for high c^H, the "specialization" equilibrium (L, H) never exists (in the presence of information frictions) because if firm L adds product H it does not give up any rents on product L.

We remain agnostic as to which equilibrium firms are more likely to play when multiple equilibria exist (i.e., for parameter values $c^H < \theta^L q^H$ and $\mu \geq \mu^*$). Still, we want to emphasize that there are theoretical reasons, beyond their empirical relevance, to believe that the "overlapping" equilibrium is compelling. First, the Pareto criterion does not allow to rule out the "overlapping" equilibrium in general, despite the fact that this criterion results in lower prices. In particular, the firm that carries product H at the "specialization" equilibrium is not necessarily better off than at the "overlapping" equilibrium, as at the former it fails to capture the profits from serving the uninformed low types. Furthermore, some authors have documented path dependency in equilibrium choices (Romero, 2015). In our setting, such path dependency suggests that the existence of the "overlapping" equilibrium for all $\mu < 1$ (in contrast to the "specialization" equilibrium, which exists only for $\mu \geq \mu^*$), together with low μ as an initial condition, might create inertia at (LH, LH) all the way down to $\mu \to 1$.[44] Last but not least, the equilibrium with overlapping qualities naturally converges to the Bertrand equilibrium as information frictions vanish, while the same is not true for the specialization equilibrium.

6.3 Comparative Statics

Combining the results of Propositions 1, 3, and 7, our final lemma evaluates how the mass of uninformed consumers affects expected market prices and expected consumer surplus at the SPE product choices. Results are illustrated in Figures 3 and 4.

[44] Consider for instance a simple repetition of our two-stage game and allow for information frictions to gradually fall. If, as initial condition, there are strong information frictions so that the "overlapping" equilibrium is unique, hysteresis would lead firms to keep on playing the same equilibrium even if the reduction in information frictions implies that the "specialization" equilibrium eventually arises. The same would apply if the costs of quality are initially high and declining. In contrast, if there are initially no search costs and the costs of quality provision are low, the market could remain at the "specialization" equilibrium as either search costs or quality costs go up, eventually giving rise to the "overlapping" equilibrium when the "specialization" equilibrium ceases to exist. However, given the overall current trend towards lower search costs, this possibility does not seem empirically relevant.

Lemma 2 *The comparative statics of expected prices and expected consumer surplus with respect to* μ *are as follows:*

(i) At the "overlapping" equilibrium, expected prices monotonically decrease in $\mu \in [0, 1]$. *Similarly, expected consumer surplus increases in* $\mu \in [0, 1]$.

(ii) If, in case of multiple equilibria, firms always play the "specialization" equilibrium, expected prices jump upward at $\mu = \mu^*$. *Similarly, expected consumer surplus jumps downward at* $\mu = \mu^*$.

Proof. See the appendix. ■

The conventional wisdom that milder information frictions lead to lower prices applies in this model only when the reduction in information frictions does not change equilibrium product lines.[45] Indeed, when lower information frictions induce firms to switch from the "overlapping" to the "specialization" equilibrium (i.e., at $\mu = \mu^*$), expected prices jump upward as firms manage to mitigate competition by differentiating their product choices. A similar result is in Kuksov (2014), who finds, in a model of horizontal differentiation, that lower search costs can give rise to greater product differentiation and hence higher prices.

Similarly, as information frictions go down, consumer surplus goes down with a discontinuity when firms switch from the "overlapping" to the "specialization" equilibrium. The discontinuity in consumer surplus is more pronounced than the discontinuity in expected prices because not only do expected prices jump upward, but also gross consumer surplus jumps downward because of incomplete price discrimination at the "specialization" equilibrium (meaning that some high types fail to buy their preferred good, while some low types fail to consume at all).

7 Extensions and Variations

In the preceding sections we characterized product and price choices in a model (i) with two possible quality levels and two consumer types, in which (ii) search cannot be conditioned on product choices (as these were assumed to be non-observable prior to search),

[45] In general, search costs are thought to relax competition, thus leading to higher retail prices, although not as intensively as the Diamond paradox would have anticipated (Diamond, 1971). There are some exceptions to this general prediction. Some recent papers have shown that search costs can lead to lower retail prices, particularly so when search costs affect the types of consumers who search. For instance, see Janssen and Shelegia (2015), Moraga-González *et al.* (2017) and Fabra and Reguant (2020).

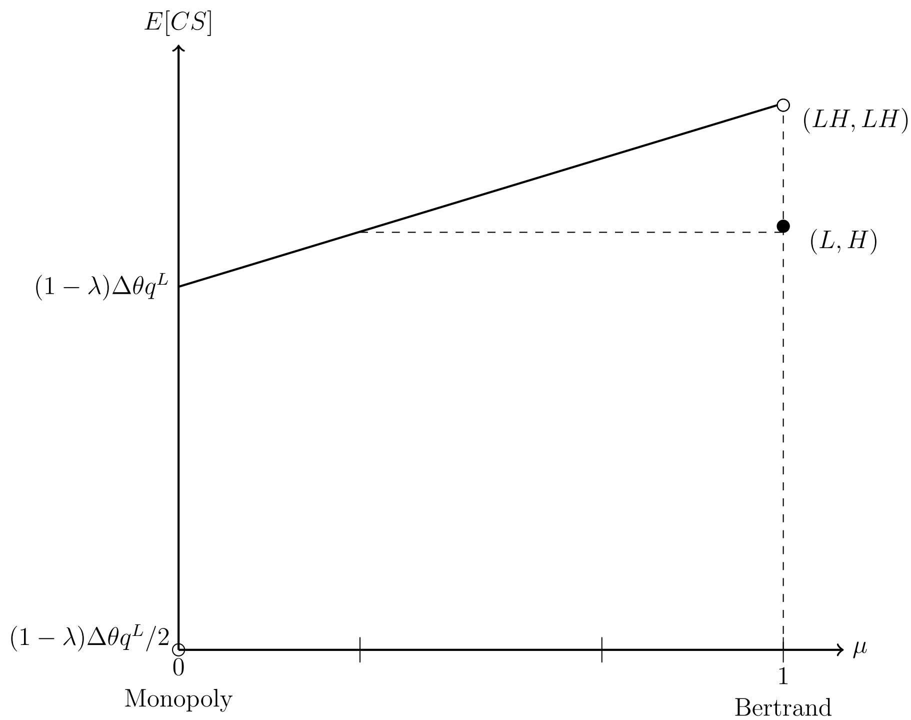

Figure 3: Expected consumer surplus as a function of μ at the SPE product choices for $c^H > \theta^L q^H$.

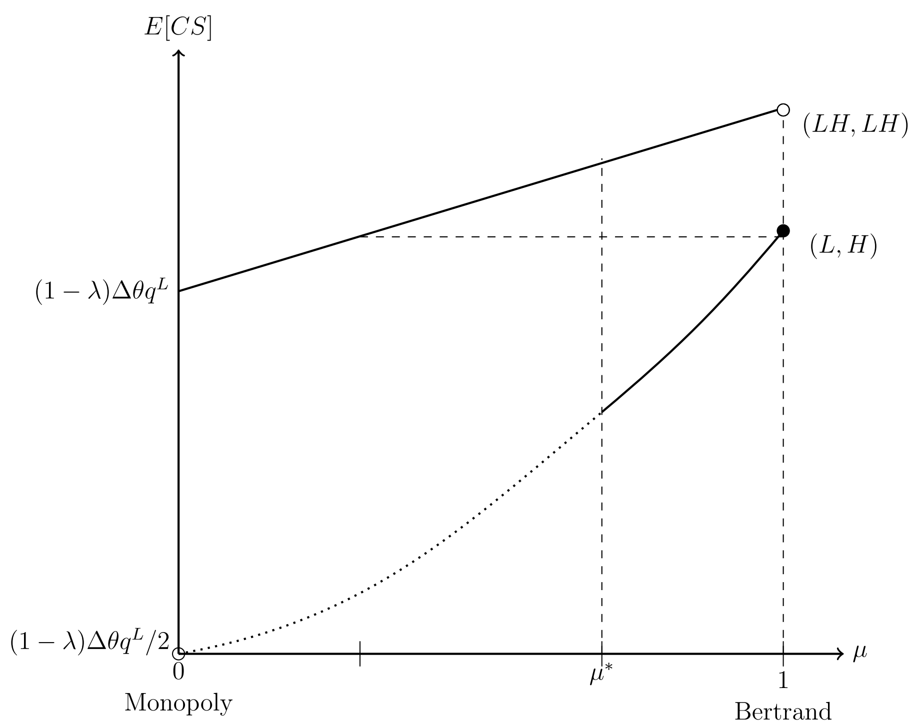

Figure 4: Expected consumer surplus as a function of μ at the SPE product choices for $c^H < \theta^L q^H$.

and in which (iii) consumers' information frictions and quality preferences are uncorrelated. In this section, we discuss how one can relax these assumptions while preserving our main results. Our focus is on the existence of the "overlapping" equilibrium.

Observable product choices and directed search by uninformed consumers In the main model, we assumed that consumers do not observe product lines prior to visiting the stores. In particular, we assumed that uninformed consumers visit one of the two stores with equal probability, regardless of their product choices. Suppose now, instead, that uninformed consumers visit the store that gives them higher expected utility, given firms' (observable) product choices and expected prices (in case of indifference, uninformed consumers visit the store that carries their preferred product).[46] Allowing search to be conditioned on product choices would strengthen our main result: when directed search is allowed, carrying multiple products would allow firms to not only better discriminate consumers, but also attract more uninformed consumers.

Directed search by uninformed consumers only affects pricing when firms have chosen asymmetric product lines (with symmetric product lines, expected prices are also symmetric, so whether search is directed or random is irrelevant). Let us consider subgame (L, LH). Now, all uninformed high types visit the multi-product firm, given that (i) the expected utility of buying product L is the same across the two stores, and (ii) at store LH, customers are indifferent between buying L or H. In turn, prices for product L must be such that uninformed low types are indifferent between visiting one store or the other (otherwise, all would visit the one charging lower prices, but this cannot constitute an equilibrium as the high-priced firm would make no sales). From our previous analysis, we know that, with an even split of uninformed consumers between the two stores, the multi-product firm charges lower prices. Hence, to rebalance the firms' pricing incentives, more than half of the uninformed low types must visit store LH until the firms' expected prices converge. Thus, since the market share of the single-product firm is lower, it makes lower profits than when product lines are non-observable, as we had assumed in the main model. In turn, this result implies that firms have no incentives to deviate from (LH, LH) to (L, LH) - their incentives to deviate are weaker than in the main model, under which (LH, LH) already constituted an equilibrium for all $\mu < 1$ (Proposition 7). Similar reasoning applies to subgame (H, LH).

In sum, our main conclusion —that the "overlapping" equilibrium is robust for all $\mu < 1$— remains valid regardless of whether product lines are observable (and there is directed search by uninformed consumers) or not. The conclusion that multi-product

[46] This interpretation of uninformed consumers as sophisticated buyers approximates that in the clearing-house model *à la* Baye and Morgan (2001).

firms tend to charge lower expected prices would have to be qualified, as with directed search, firms are expected to charge the same prices even though multi-product firms make higher profits by attracting more customers.

Correlation between information frictions and quality preferences Finally, we have assumed so far that informed and uninformed consumers are equally likely to be either high or low types. However, this assumption may not hold in practice. For instance, if low types are lower income consumers with more time to search, then uninformed consumers are more likely to be high types. Alternatively, if high types enjoy shopping for their preferred (high quality) product, then uninformed consumers are more likely to be low types. Ultimately, this is an empirical question whose answer may vary depending on the product type or context considered. However, as far as the predictions of the model are concerned, it is inconsequential whether the correlation between information frictions and quality preferences is positive, negative, or non-existent.

To formalize this result, one can introduce parameters λ^I and λ^{NI}, representing the fraction of low types among informed and uninformed consumers, respectively, i.e., $\lambda^I\mu + \lambda^{NI}(1-\mu) = \lambda$. If $\lambda^I > \lambda^{NI}$, there is positive correlation between information frictions and quality types, as the fraction of low types is higher among informed consumers than among uninformed consumers. In line with the main text, we assume that this correlation is not too strong, so that monopoly profits from uninformed are still maximized by selling both products, as in assumption $(A1)$.

The analysis of product and price choices without information frictions remains intact since all consumers are informed by definition. As for the analysis with information frictions, profits on good H are proportional to $\left(1-\lambda^{NI}\right)$ and those on good L are proportional to λ^{NI}, implying that the incentive structure remains unchanged. As such, the "overlapping" equilibrium always exists, just as in the case with no correlation between search and quality preferences.

8 Conclusions

In this paper, we have analyzed the impact of information frictions on quality choices followed by price competition. We have found that the equilibrium in which firms carry overlapping product lines always exists, with or without information frictions. In contrast, we have also shown that the equilibrium with non-overlapping quality choices, as originally proposed in Champsaur and Rochet (CR)'s (1989) influential paper, is not particularly robust: it exists only under mild information frictions if the costs of providing high quality are sufficiently small. In particular, if providing high quality is particularly costly, CR's equilibrium fails to exist even when the mass of uninformed consumers is infinitesimally small. This finding casts doubts on the prediction that strategic incentives alone induce firms to soften competition by carrying non-overlapping product lines. Our results extend to more general settings, including cases of more than two goods and consumer types, more than two firms, directed search by uninformed consumers, and the possibility that information frictions and quality tastes are positively or negatively correlated.

We have shown that, through product choice, information frictions can have important implications for market outcomes beyond their well studied price effects. In particular, we have shown that analyzing the price effects of information frictions without endogenizing product choices can sometimes lead to overestimating their anticompetitive effects. This is the case, for instance, when the addition of information frictions induces firms to carry overlapping products, creating head-to-head competition.

The multi-product nature of firms also adds important twists to the analysis of competition in the presence of information frictions. An important departure from Varian (1980) is that goods within a store cannot be priced independently from each other. In particular, the incentives to separate both consumer types impose an upper (lower) bound on the highest price that can be charged for a high (low) quality good, given the price of the low (high) quality one. This holds true even for a single-product firm competing with a multi-product one, since through competition, the effects of price discrimination by the multi-product firm affect pricing by the single-product firm. In line with Varian (1980), we have also shown that information frictions give rise to price dispersion when the two competing firms carry multiple products—a possibility not considered by Varian (1980).

Admittedly, there are several determinants of firms' product choices beyond those studied in this paper. In particular, throughout the analysis we have assumed that firms do not incur any fixed cost of carrying a product. This modelling choice was meant to highlight the strategic motives underlying product choice. However, fixed costs of carrying a product (which could arguably be higher for high-quality products)[47] could induce firms to offer fewer, possibly non-overlapping products. Our prediction is not that competitors should always carry overlapping product lines. Rather, our analysis

[47] In some cases, such costs can be substantial, e.g., firms have to advertise that they are carrying an additional product, or the transaction costs of dealing with an additional provider can sometimes be high. The marketing literature has analyzed several factors explaining the limited number of products sold per firm. For instance, Villas-Boas (2004) analyzes product line decisions when firms face costs of communicating about the different products they carry to their customers. They show that costly advertising can induce firms to carry fewer products as well as to charge lower prices for their high-quality goods.

suggests that if their product lines do not overlap in markets with information frictions, the reasons must be other than the firms' attempts to soften competition through product choice—for instance, due to fixed costs.

Appendix A: Selected Proofs

Proof of Proposition 1 [SPE under $\mu = 1$] We show that the specialization equilibrium (L, H) constitutes a SPE. First, at subgames (LH, LH), (L, L) and (H, H), both firms make zero profits. Second, at subgame (L, LH) the low-quality product is priced at marginal cost c^L while the high-quality product is sold at the highest price that satisfies the high types' incentive compatibility constraint, i.e., $c^L + \theta^H \Delta q$. Firm L makes zero profits while firm LH gets a payoff of $(1-\lambda)\varphi^H$, which equals its minimax. Third, at subgame (H, LH), the high-quality product is priced at marginal cost c^H while the low-quality product is sold at the highest price that satisfies the low types' incentive compatibility constraint and participation constraints, i.e., $\min\left\{c^H - \theta^L \Delta q, \theta^L q^L\right\}$. Firm H makes zero profits while firm LH makes profits $\lambda\pi^L$ if $c^H > \theta^L q^H$ or $\lambda\varphi^L$ otherwise, i.e., its minmax. Finally, at subgame (L, H) the equilibrium is in mixed strategies. For the purposes of this proof, it suffices to put bounds on equilibrium profits. Minmax profits for each firm are computed by characterizing the firm's best response to the rival pricing its good at marginal cost. Following our previous analysis, the minmax profits for the H firm are $(1-\lambda)\varphi^H > 0$, while the minmax profits for the L firm are $\lambda\pi^L > 0$ if $c^H > \theta^L q^H$ or $\lambda\varphi^L > 0$ otherwise. Since at the mixed strategy equilibrium firms always price above marginal costs (otherwise they would have zero profits, but this cannot be since their minmax profits are positive), equilibrium profits are *strictly above* the minimax whenever the participation constraint is not binding. The only case where above marginal cost pricing does not necessarily imply that firm L's profits are strictly above its minmax is when $c^H > \theta^L q^H$, as in this case firm L's best response is the the same regardless of whether firm H prices at c^H or above.[48] Indeed, for the case $c^H > \theta^L q^H$, equilibrium profits are exactly equal to the minmax $\lambda\pi^L$. To see this, note that at the mixed-strategy equilibrium (MSE), the upper bounds of firms' price supports are the

[48] It is straightforward to see that in a mixed strategy equilibrium we must have $\underline{p}^H > c^H$ and $\underline{p}^L > c^L$; otherwise, each firm's profits would be zero, but this leads to a contradiction since profits cannot be below the minimax. Hence, firm H would never like to price lower than $\underline{p}^L + \theta^H \Delta q > c^L + \theta^H \Delta q > c^H$. Since at a price $\underline{p}^L + \theta^H \Delta q$ firm H would at least be serving high types, its profits must be strictly greater than its minmax $(1-\lambda)\varphi^H$. Similarly, if $\underline{p}^H > \theta^L q^H$, firm L would be a monopolist over low types, so it could always secure profits of at least $\lambda\pi^L$. If $\underline{p}^H < \theta^L q^H$, firm L would never like to charge prices lower than $\underline{p}^H - \theta^L \Delta q > c^H - \theta^L \Delta q$. Since at a price $\underline{p}^H - \theta^L \Delta q$ firm L would at least be serving low types, its profits must be strictly greater than its minmax $\lambda\varphi^L$.

constrained monopoly prices. Furthermore, firm L has to play a probability mass at its upper bound. Otherwise, firm H would make zero profits at its upper bound (as all consumers would strictly prefer to buy from firm L), but this cannot be the case since its minmax is strictly positive. Finally, the two firms cannot put positive mass at their upper bounds as firm L would be better off putting all its mass slightly below its upper bound (so as to attract all consumers whenever the rival plays the mass at the upper bound). It thus follows that when firm L plays its upper bound, the rival is pricing below its upper bound with probability one. Hence, at the upper bound firm L serves only low types, thus making profits that exactly equal its minmax, $\lambda\pi^L$.

We are now ready to show that (L, H) is a SPE. Starting at (L, H), firm H does not want to carry good L as at (L, LH) its profits are equal to the minmax, while they are strictly above that level at (L, H). Similarly, firm L does not want to carry good H as at (LH, H) its profits are equal to the minmax, while at (L, H) its profits are (weakly) greater than its minmax.

Last, we characterize the MSE. Suppose that the rival chooses L with probability α and H with probability $(1-\alpha)$. Equating the profits from choosing L and H,

$$\alpha\Pi(L, L) + (1-\alpha)\Pi(L, H) = \alpha\Pi(H, L) + (1-\alpha)\Pi(H, H)$$

Since $\Pi(H, H) = \Pi(L, L) = 0$, solving for α,

$$\alpha = \frac{\Pi(L, H)}{\Pi(L, H) + \Pi(H, L)}.$$

Thus implying that equilibrium profits at the MSE are

$$\frac{\Pi(L, H)\Pi(H, L)}{\Pi(L, H) + \Pi(H, L)}.$$

This equilibrium constitutes a SPE if and only if it is not dominated to choosing LH, i.e.,

$$\Pi(L, H)[\Pi(L, L) - \Pi(LH, L)] + \Pi(H, L)[\Pi(L, H) - \Pi(LH, H)] \geq 0$$

The first term is negative, while the sign of the second term depends on c^H : (i) if $c^H \geq \theta^L q^H$, it is negative, implying that LH dominates the MSE candidate, which therefore does not exist; on the contrary, (ii) if $c^H < \theta^L q^H$, the second term is positive, implying that a MSE cannot be ruled out, particularly for low values of c^H, which is when the second term is higher (note that as $c^H \to \theta^L q^H$ the second term is close to zero, so the MSE is ruled out for some $c^H < \theta^L q^H$). Therefore, for those parameter values for which the above inequality holds, a MSE exists. *Q.E.D.*

Proof of Proposition 2 [SPE under $\mu \to 1$] The results on existence and uniqueness of the "overlapping" equilibrium (LH, LH) under the assumption $\mu \to 1$ are a particular case of the proof of Proposition 7. The proof of non-existence of the "specialization" equilibrium (L, H) for $\mu \to 1$, for the case $c^H \geq \theta^L q^H$ is also contained in the proof of Proposition 7. Hence, it only remains to prove that $c^H < \theta^L q^H$ implies the existence of the "specialization" equilibrium. The proof of Proposition 1 above shows that $\Pi(L, H)$ and $\Pi(H, L)$ are *strictly* above the minmax, while the proofs of Proposition 5 and 4 show that $\Pi(LH, H)$ and $\Pi(LH, L)$ are equal to their minmax as $\mu \to 1$. Deviating to $\Pi(L, L)$ or $\Pi(H, H)$ is also not profitable given that firms would make almost zero profits. It follows that no firm wants to deviate from (L, H). *Q.E.D.*

Proof of Lemma 1 Argue by contradiction and suppose first that a firm chooses $\Delta p > \theta^H \Delta q$, which together with $(A2)$ implies that

$$\Delta p > \theta^H \Delta q > \Delta c$$

and that nobody is buying the high-quality product. But if the firm reduces p^H so that $\theta^H \Delta q > \Delta p > \Delta c$, it would not affect low types' decisions given that by $(A2)$, $\theta^L \Delta q < \Delta c$ implies $\theta^L \Delta q < \Delta p$ which can be re-written as $\theta^L q^L - p^L > \theta^L q^H - p^H$. Instead, high-type consumers would be better off (note that $\theta^H \Delta q > \Delta p$ can be re-written as $\theta^H q^H - p^H > \theta^H q^L - p^L$), some strictly, the firm's uninformed consumers, while others weakly, the informed consumers. Such a price reduction turns out to be a profitable deviation since $\Delta p > \Delta c$ implies $p_H - c_H > p_L - c_L$.

Suppose now that $\Delta p < \theta^L \Delta q$, which together with $(A2)$ implies that

$$\Delta c > \theta^L \Delta q > \Delta p$$

and that nobody is buying the low quality product. But if the firm reduces p^L so that $\Delta c > \Delta p > \theta^L \Delta q$, it would not affect high types' decisions given that by $(A2)$, $\theta^H \Delta q > \Delta c$ implies $\theta^H \Delta q > \Delta p$ which can be re-written as $\theta^H q^H - p^H > \theta^H q^L - p^L$. Instead, low-type consumers would be better off (note that $\Delta p > \theta^L \Delta q$ can be re-arranged as $\theta^L q^L - p^L > \theta^L q^H - p^H$), some strictly, the firm's uninformed consumers, while others weakly, the informed consumers. Again, such a price reduction turns out to be a profitable deviation since $\Delta c > \Delta p$ implies $p_L - c_L > p_H - c_H$. *Q.E.D.*

Proof of Proposition 3 [pricing at subgame (LH, LH)] To prove that a pure-strategy equilibrium does not exist we proceed by contradiction. There are two cases to consider. The first one is when a firm, say firm 1, is serving only uninformed consumers. If so, firm 1 must be charging the constrained monopoly prices (i.e., $p_1^L = \theta^L q^L$ and

$p_1^H = \theta^L q^L + \theta^H \Delta q$). But if so, its rival's best response is to price both goods slightly below such prices, to which firm 1 would in turn respond with a slight undercut of its own, and so on.

The second case is when firm 1 is also serving some informed consumers, along its uninformed consumers. There are three possibilities here: serving all informed consumers, serving only informed high types, and serving only informed low types. If firm 1 is serving all informed consumers, firm 2 must be serving only its uninformed consumers, and we are back to the previous case. Consider then the second possibility that firm 1 is serving only informed high types (while firm 2 serves informed low types), along its uninformed consumers. Firms' prices must be incentive compatible (Lemma 1) and satisfy *(a)* $p_1^H < p_2^H$ and *(b)* $p_1^L > p_2^L$. Furthermore, a firm's prices must constitute a best response to their rival's subject to *(a)* and *(b)*, and vice versa. In particular, this implies that $p_2^L = p_1^L - \epsilon$ and $p_1^H = p_2^H - \epsilon$, with $\epsilon \to 0$. But if so, both firms would have incentives to deviate. Firm 1 would deviate to $p_1^L = p_2^L - \epsilon$ so as to capture informed low types as well, and firm 2 would deviate to $p_2^H = p_1^H - \epsilon$ so as to capture informed high types as well. Similar arguments allow us to discard the remaining possibility that firm 1 is serving only informed low-type consumers (and firm 2 informed high-type consumers), along its uninformed consumers. Thus, the equilibrium must be in mixed-strategies.

For the rest of the proof we focus on characterizing symmetric mixed strategy equilibria. Standard arguments imply that there are no holes in the support and that firms play no mass point at any price of the support, including the upper bound (see, for instance, Narasimhan, 1998). (i) At the upper bound, firms serve the uninformed consumers only. Since profits are increasing in prices subject to $\left(IC^H\right)$, the optimal prices at the upper bounds are $\overline{p}^H = \theta^H q^H - q^L \Delta\theta$ and $\overline{p}^L = \theta^L q^L$, so that $\Delta\overline{p} = \theta^H \Delta q$.

We now demonstrate (ii), i.e., that at the lower bound $\Delta\underline{p} < \theta^H \Delta q$. Suppose otherwise that the price gap $p^H - p^L$ is constant and equal to $\theta^H \Delta q$ at and in the neighborhood of the lower bound (or throughout the entire price support for that matter). When a firm plays $(\underline{p}^H, \underline{p}^L)$ it obtains

$$\Pi(\underline{p}^H, \underline{p}^L; LH, LH) = \left(\mu + \frac{1-\mu}{2}\right) \lambda(\underline{p}^L - c^L) + \left(\mu + \frac{1-\mu}{2}\right)(1-\lambda)(\underline{p}^H - c^H).$$

Using

$$\begin{aligned} \Pi(\underline{p}^H, \underline{p}^L; \cdot) = \bar{\pi} &\equiv (1-\mu)\left[\lambda\pi^L + (1-\lambda)\left(\pi^H - \Delta\theta q^L\right)\right]/2 \\ &= (1-\mu)\left[\pi^L + (1-\lambda)\left(\theta^H \Delta q - \Delta c\right)\right]/2, \end{aligned}$$

the payoff at the upper bound, and the assumption that $\underline{p}^H - \underline{p}^L = \theta^H \Delta q$ we obtain

$$\underline{p}^H - c^H = \frac{1-\mu}{1+\mu}(\bar{p}^H - c^H) + \lambda \frac{2\mu}{1+\mu}\varphi^H \tag{6}$$

and

$$\underline{p}^L - c^L = \frac{1-\mu}{1+\mu}(\bar{p}^L - c^L) - (1-\lambda)\frac{2\mu}{1+\mu}\varphi^H. \tag{7}$$

We now compute the marginal distribution function $F^i(p^i)$ (since Δp is fixed in the neighborhood of the lower bound there is just one distribution to consider, say $F(p^i)$). First, notice that if one firm plays something in the support, the other firm never wants to deviate and serve just high types at a price $\theta^H q^H$, because according to $(A1)$ the payoff of doing so would be strictly lower. Thus, to obtain the cdf $F(p^H)$ around the lower bound, notice that playing any pair p^H and $p^L = p^H - \theta^H \Delta q$ around the lower bound yields an expected payoff of

$$\begin{aligned}\Pi(p^H, p^L; \cdot) &= (1-\lambda)(p^H - c^H)\left[\frac{1-\mu}{2} + \mu(1 - F(p^H))\right] + \\ &\quad \lambda(p^L - c^L)\left[\frac{1-\mu}{2} + \mu(1 - F(p^H))\right]\end{aligned}$$

where $1 - F(p^H) = 1 - F(p^L = p^H - \theta^H \Delta q)$ is the probability to attract all informed consumers, high and low types. Rearranging terms and using $\Pi(p^H, p^L = p^H - \theta^H \Delta q) = \bar{\pi}$ leads to

$$\frac{1-\mu}{2}(\bar{p}^H - c^H) = \left[1 - F(p^H)\right]\left[\mu(p^H - c^H) - \lambda\mu\varphi^H\right] \tag{8}$$

and

$$\frac{1-\mu}{2}\left(\bar{p}^L - c^L\right) = [1 - F(p^L)[\mu(p^L - c^L) + (1-\lambda)\mu\varphi^H]. \tag{9}$$

Evaluating $F(\underline{p}^H) = F(\underline{p}^L) = 0$ in (8) and (9) yields

$$\underline{p}^H - c^H = \frac{1-\mu}{2\mu}(\bar{p}^H - c^H) + \lambda\varphi^H$$

and

$$\underline{p}^L - c^L = \frac{1-\mu}{2\mu}(\bar{p}^L - c^L) - (1-\lambda)\varphi^H$$

which, since $\mu < 1$, are greater than (6) and (7), respectively; a contradiction.

Proof for the third item in the proposition follows directly from Lemma 1. *Q.E.D.*

Proof of Proposition 7 [quality choices] Each firm has four potential choices: $\{\varnothing, L, H, LH\}$. On the one hand, to prove that (LH, LH) is a SPE of the game for all $\mu < 1$, just note that all equilibrium payoffs $\Pi(LH, LH)$, $\Pi(H, LH)$ and $\Pi(L, LH)$ are proportional to $(1-\mu)/2$ so that $(A1)$ allows to conclude that $\Pi(LH, LH)$ is the greatest among these, just as in the monopoly case.

On the other hand, to find the conditions under which (L, H) is an equilibrium, we need to assess firm L's deviation gains when also carrying good H (it is easy to check that this is the critical deviation; for instance, let $\mu \to 1$ and use $(A2)$ to note that firm L's deviation gains are greater than firm H's, i.e., $\Pi(LH,H) - \Pi(L,H) > \Pi(LH,L) - \Pi(H,L)$). Firm L's deviation gain is equal to

$$\Pi(L,H) - \Pi(LH,H) = \frac{1+\mu}{2}\lambda\left(\pi^L - \varphi^L\right) - \frac{1-\mu}{2}(1-\lambda)\left(\pi^H - \pi^L\right) \qquad (10)$$

(See the proofs of Propositions 5 and 6 in the online appendix for the relevant payoffs). Solving for μ, the above profit difference is positive iff $\mu \geq \mu^*$, where

$$\mu^* = \frac{(1-\lambda)\left(\pi^H - \pi^L\right) - \lambda\left(\pi^L - \varphi^L\right)}{(1-\lambda)\left(\pi^H - \pi^L\right) + \lambda\left(\pi^L - \varphi^L\right)}.$$

Note that when $c^H \geq \theta^L q^H$, $\pi^L = \varphi^L$ and $\mu^* = 1$, making $c^H \geq \theta^L q^H$ a sufficient condition for the uniqueness of (LH, LH). Furthermore, taking the derivative of μ^* with respect to c^H shows that

$$\frac{\partial \mu^*}{\partial c^H} = -2\lambda(1-\lambda)\frac{\left(\pi^H - \pi^L\right) - \left(\pi^L - \varphi^L\right)}{(1-\lambda)\left(\pi^H - \pi^L\right) + \lambda\left(\pi^L - \varphi^L\right))^2}.$$

So that

$$\begin{aligned} sign\left\{\frac{\partial \mu^*}{\partial c^H}\right\} &= -sign\left\{\left(\pi^H - \pi^L\right) - \left(\pi^L - \varphi^L\right)\right\} \\ &= -sign\left\{c^L - \theta^L q^H\right\} > 0. \end{aligned}$$

Finally, there might also exist a symmetric MSE such that firms choose L and H randomly, just as shown in the proof of Proposition 2. This equilibrium constitutes a SPE if and only if it is not dominated to choosing LH, i.e.,

$$\Pi(L,H)\left[\Pi(L,L) - \Pi(LH,L)\right] + \Pi(H,L)\left[\Pi(L,H) - \Pi(LH,H)\right] \geq 0$$

or equivalently, iff

$$\Pi(L,H) - \Pi(LH,H) \geq \frac{\Pi(L,H)}{\Pi(H,L)}\left[\Pi(LH,L) - \Pi(L,L)\right] > 0.$$

Hence, whereas the existence of the asymmetric PSE (L,H) requires the profit difference (10) to be positive, the existence of the MSE requires such a difference to be greater than a strictly positive number. If we denote with μ^{**} the critical value for the existence of the MSE, we must then have $\mu^{**} \geq \mu^*$. It thus follows that for $\mu < \mu^*$, the unique equilibrium (either pure or mixed) is (LH, LH). *Q.E.D.*

Proof of Lemma 2 [prices and consumers surplus at the SPE] It is straightforward to see that, conditional on firms playing (LH, LH), expected prices are decreasing in μ. Since there is full discrimination, total consumption of each good remains fixed so that total surplus is given by $\lambda\pi^L + (1-\lambda)\pi^H$, irrespectively of μ. Since profits in equation (3) decrease in μ, consumer surplus must increase in μ. In turn, this implies that expected prices must be decreasing in μ. For given parameter values, competition is stronger at subgame (LH, LH) than at (L, H). Hence, expected prices at the former must be lower and consumer surplus must be higher. Thus, as μ goes down, expected prices (consumer surplus) at (LH, LH) decrease continuously until they jump up when firms start playing (L, H), either at $\mu \to 1$ or at $\mu \to \mu^*$ depending on equilibrium selection. Similarly, as μ goes down, consumer surplus at (LH, LH) increases continuously until it jumps down when firms start playing (L, H), either at $\mu \to 1$ or at $\mu \to \mu^*$ depending on equilibrium selection. *Q.E.D.*

Appendix B: Additional Results

Mixed strategy pricing equilibrium at (LH, LH) Consider pricing at the subgame (LH, LH). As we argued in Section 3, there are potentially multiple mixed strategy, outcome-equivalent, equilibria. Because the incentive compatibility constraint of the high types is binding at the monopoly solution, a natural equilibrium to consider is one in which firms keep on pricing the low-quality product as if they were just selling that product, but adjust their pricing for the high-quality one. The following Lemma characterizes such an equilibrium.

Lemma 3 *Given product choices (LH, LH), there exists a mixed-strategy equilibrium in which firms choose p^L in $\left[\underline{p}^L, \overline{p}^L\right]$ according to the (conditional and marginal) distribution function*

$$F^L(p^L) = \frac{1+\mu}{2\mu} - \frac{1-\mu}{2\mu}\frac{(\overline{p}^L - c^L)}{(p^L - c^L)}$$

and such that, for given p^L, the price p^H is chosen in $\left[\underline{p}^H, \overline{p}^H\right]$ to satisfy

$$\frac{p^H - c^H}{p^L - c^L} = \frac{\overline{p}^H - c^H}{\overline{p}^L - c^L} \tag{11}$$

where

$$\underline{p}^i = c^i + \frac{1-\mu}{1+\mu}(\overline{p}^i - c^i) > c^i,$$

and $\overline{p}^i$ are the (constrained) monopoly prices, for $i = L, H$.

Proof of Lemma 3: We want to show that the equilibrium in the statement of the lemma is indeed an equilibrium. First, firms could deviate by playing the price pairs in the support with different probabilities, while still choosing price pairs that satisfy incentive compatibility. However, this is unprofitable given that all price-pairs in the support give equal expected profits. Indeed, the equilibrium has been constructed so that

$$(p^L - c^L)\left[\frac{1-\mu}{2} + \mu(1 - F^L(p^L))\right] = \frac{1-\mu}{2}(\overline{p}^L - c^L) = \frac{1+\mu}{2}(\underline{p}^L - c^L)$$

and

$$(p^H - c^H)\left[\frac{1-\mu}{2} + \mu(1 - F^H(p^H))\right] = \frac{1-\mu}{2}(\overline{p}^H - c^H) = \frac{1+\mu}{2}(\underline{p}^H - c^H)$$

with the ratio (11) derived in order for the price pair $\left(p^H, p^L\right)$ to satisfy $F^H(p^H) = F^L(p^L)$, i.e., the choice of p^L results in a choice of p^H, so that the prices satisfying that ratio are played with equal probability. Therefore, expected profits at the proposed equilibrium are as in (3).

Second, firms could deviate by choosing p^L and p^H not satisfying equation (11) while still satisfying incentive compatibility. Again, these deviations are not profitable since all the prices in the support give equal profits. Deviating to prices that do not satisfy incentive compatibility is unprofitable because of Lemma 1.

Finally, firms could deviate by playing price pairs outside the support. Choosing any prices above $\left(\overline{p}^L, \overline{p}^H\right)$ as defined above is unprofitable, as at these prices the firm is only selling to the uninformed consumers and $\left(\overline{p}^L, \overline{p}^H\right)$ are the optimal monopoly prices. Choosing any prices below $\left(\underline{p}^L, \underline{p}^H\right)$ as defined above is unprofitable, as at these prices the firm is inelastically selling to all consumers with probability one and would thus gain by rasing the price up to $\left(\underline{p}^L, \underline{p}^H\right)$. *Q.E.D.*

The proposed equilibrium has several appealing features. While firms price the low-quality product as if they were just selling that product (as in Varian's model), on average they choose lower prices for the high-quality product than when they only sell that product. This is a direct implication of the fact that the firm cannot extract all the surplus of the uninformed high types. Indeed, the resulting distribution for p^H,

$$F^H(p^H) = \frac{1+\mu}{2\mu} - \frac{1-\mu}{2\mu}\frac{(\overline{p}^H - c^L)}{(p^H - c^L)}$$

has the same functional form as in Varian. However, since the upper bound $\overline{p}^H$ is the constrained monopoly price, the whole distribution puts higher weight on lower prices all along the support than in the independent products case.

Under this equilibrium, the choice of p^L results in a unique choice of p^H such that the relative profit margin of the two products remains constant along the whole support; see equation (11).[49] In particular, the relative markups of the two products are the same as under monopoly. That is, under this equilibrium, competition affects the *price levels* but not the *price structure* within the firm.[50] The reason why the *relative profitability of two products is kept constant* simply derives from choosing p^H so as to make the incentive compatibility constraint binding, for every p^L.

The price difference that is embodied in this price structure can be expressed as

$$\Delta p = \kappa \theta^H \Delta q + (1 - \kappa) \Delta c.$$

Consistently with Lemma 1, the price difference is a weighted average between $\theta^H \Delta q$ (i.e., the price difference at the monopoly solution) and Δc (i.e., the price difference at the competitive solution), where the weight $\kappa = \left(p^L - c^L\right) / \left(\overline{p}^L - c^L\right)$ represents the distance to the upper bound. At the upper bound, when the incentive compatibility constraint of the high types is binding, the price difference is maximal, $\Delta \overline{p} = \theta^H \Delta q$. As we move down the support, the incentive compatibility constraint is satisfied with slack and the price difference narrows down. The difference is minimal at the lower bound, when $\kappa = (1 - \mu) / (1 + \mu)$. Importantly, as μ approaches one, the prices at the lower bound converge to marginal costs, and the price gap approaches Δc. The equilibrium would thus collapse to the competitive solution. On the other extreme, as μ approaches zero, the prices at the lower bound converge to monopoly prices so that the price gap approaches $\theta^H \Delta q$. The equilibrium would thus collapse to the monopoly solution.

References

[1] Armstrong, M. (2017), Ordered Consumer Search, *Journal of the European Economic Association* 15(5), 989-1024.

[2] Anderson, S., and Renault, R. (1999), Pricing, Product Diversity, and Search Costs: A Bertrand-Chamberlin-Diamond Model, *RAND Journal of Economics* 30(4), 719-735.

[49] Clearly, there exists another equilibrium with the same price supports and the same price distribution for good L but in which the firm randomizes the price of good H, given the choice of p^L, such that the two prices remain incentive compatible. Again, this multiplicity is inconsequential for the purposes of this analysis as all equilibria yield equal expected profits.

[50] Note that in this equilibrium, the prices of the two products within a firm are positively correlated. This is in contrast to what the literature on multi-product loss-leading concludes. However, as discussed in the introduction, that literature applies to setups in which goods are complements and consumers buy more than one- in contrast to the assumptions made in this paper.

[3] Baye, M., and J. Morgan (2001), Information Gatekeepers on the Internet and the Competitiveness of Homogeneous Product Markets, *American Economic Review* 91(3), 454-474.

[4] Borenstein, S., and Rose, N.L. (1994), Competition and Price Dispersion in the U.S. Airline Industry, *Journal of Political Economy* 102(4), 653-68.

[5] Brander, J., and Eaton, J. (1984), Product-line Rivalry, *American Economic Review* 74(3), 323-334.

[6] Burdett, K., and Judd, K.L. (1983), Equilibrium Price Dispersion, *Econometrica* 51(4), 955-969.

[7] Busse, M., and Rysman, M. (2005), Competition and Price Discrimination in Yellow Pages Advertising, *RAND Journal of Economics* 36(2), 378-390.

[8] Bar-Isaac, H., Caruana, G., and Cuñat, V. (2012), Search, Design, and Market Structure, *American Economic Review* 102(2), 1140-60.

[9] Chamberlin, E. H. (1933), The Theory of Monopolistic Competition. Cambridge, MA: Harvard University Press.

[10] Champsaur, P., and Rochet, J.-C. (1989), Multiproduct Duopolists, *Econometrica* 57(3), 533-557.

[11] Chen, L., Mislove, A., and Wilson, C. (2016), An Empirical Analysis of Algorithmic Pricing on Amazon Marketplace, mimeo.

[12] Crawford, G., and Shum, M. (2007), Monopoly Quality Degradation in Cable Television, *Journal of Law and Economics* 50(1), 181-219.

[13] Diamond, P. (1971), A Model of Price Adjustment, *Journal of Economic Theory* 3(2), 156-168.

[14] De Los Santos, B., Hortaçsu, A., and Wildenbeest, M. (2012), Testing Models of Consumer Search Using Data on Web Browsing and Purchasing Behavior, *American Economic Review* 102(6), 2955-80.

[15] Ellison, G. (2005), A Model of Add-on Pricing, *Quarterly Journal of Economics* 120(2), 585-637.

[16] Ershov, D. (2018), Competing with Superstars in the Mobile App Market, NET Institute working paper 18-02.

[17] Fabra, N., and Reguant, M. (2020), A Model of Search with Price Discrimination, *European Economic Review*, forthcoming.

[18] Gal-Or, E. (1983), Quality and Quantity Competition, *Bell Journal of Economics* 14(3), 590-600.

[19] Gamp, T. (2019), Guided Search, mimeo.

[20] Garrett, D., Gomes, R., and Maestri, L. (2019), Competitive Screening under Heterogeneous Information, *Review of Economic Studies* 86(4), 1590-1630.

[21] Gilbert, R., and Matutes, C. (1993), Product Line Rivalry with Brand Differentiation, *The Journal of Industrial Economics* 41(3), 223-240.

[22] Hong, H., and Shum, M. (2006), Using Price Distributions to Estimate Search Costs, *RAND Journal of Economics* 37(2), 257-75.

[23] Iyer, G., and Seetharaman, P.B. (2003), To Price Discriminate or Not: Product Choice and the Selection Bias Problem, *Quantitative Marketing and Economics* 1(2), 155-178.

[24] Janssen, M., and Moraga, J.L. (2004), Strategic Pricing, Consumer Search and the Number of Firms, *Review of Economic Studies*, 71(4), 1089-1118.

[25] Janssen, M., and S. Shelegia (2015), Consumer Search and Double Marginalization, *American Economic Review* 105(6), 1683-1710.

[26] Johnson, J.P., and Myatt, D.P. (2003), Multiproduct Quality Competition: Fighting Brands and Product Line Pruning, *American Economic Review* 93(3), 749-774.

[27] Kaplan, G., Menzio, G., Rudanko, L., and Thachter, N. (2019), Relative Price Dispersion: Evidence and Theory, *American Economic Journal: Microeconomics* 11(3), 68-124.

[28] Klemperer, P. (1992), Equilibrium Product Lines: Competing Head-to-Head May be Less Competitive, *American Economic Review* 82(4), 740-55.

[29] Kuksov, D. (2004), Buyer Search Costs and Endogenous Product Design, *Marketing Science* 23(4), 490-499.

[30] Leslie, P. (2004), Price Discrimination in Broadway Theater, *RAND Journal of Economics* 35(3), 520-541.

[31] McManus, B. (2007), Nonlinear Pricing in an Oligopoly Market: The Case of Specialty Coffee, *RAND Journal of Economics* 38(3), 513-533.

[32] McAfee, R. P. (1995), Multiproduct Equilibrium Price Dispersion, *Journal of Economic Theory*, 67(1), 83-105.

[33] Moraga González, J.L., Sandor, Z. and Wildenbeest, M. (2017), Prices and Heterogeneous Search Costs, *RAND Journal of Economics* 48(1), 125-146.

[34] Moorthy, K.S. (1988), Product and Price Competition in a Duopoly, *Marketing Science* 7(2), 141-168.

[35] Mussa, M., and Rosen, S. (1978), Monopoly and Product Quality, *Journal of Economic Theory* 18(2), 301-317.

[36] Narasimhan, C., (1988), Competitive Promotional Strategies, *Journal of Business* 61(4), 427-449.

[37] Nevo, A., and Wolfram, C. (2002), Why Do Manufacturers Issue Coupons? An Empirical Analysis of Breakfast Cereals?, *RAND Journal of Economics* 33(2), 319-339.

[38] Petrikaite, V. (2018), Consumer Obfuscation by a Multiproduct Firm, *RAND Journal of Economics* 49(1), 1-282.

[39] Rhodes, A. (2015), Multiproduct Retailing, *Review of Economic Studies* 82(1), 360-390.

[40] Rhodes, A., and Zhou, J. (2019), Consumer Search and Retail Market Structure, *Management Science* 65(6), 2607-2623.

[41] Romero, J. (2015), The Effect of Hysteresis on Equilibrium Selection in Coordination Games, *Journal of Economic Behavior and Organization* 111(3), 88-105.

[42] Seim, K. and Sinkinson, M. (2016), Mixed Pricing in Online Marketplaces, *Quantitative Marketing and Economics* 14(2), 129-155.

[43] Shaked, A., and Sutton, J. (1982), Relaxing Price Competition Through Product Differentiation, *Review of Economic Studies* 49(1), 3-13.

[44] Shelegia, S. (2012), Multiproduct Pricing in Oligopoly, *International Journal of Industrial Organization*, 30(2), 231-242.

[45] Shepard, A. (1991), Price Discrimination and Retail Configuration, *Journal of Political Economy* 99(1), 30-53.

[46] Stahl, D.O. (1989), Oligopolistic Pricing with Sequential Consumer Search, *American Economic Review* 79(4), 700-712.

[47] Stole, L. (1995), Nonlinear Pricing and Oligopoly, *Journal of Economics & Management Strategy* 4(4), 529-562.

[48] Stole, L. (2007), Price Discrimination and Competition, *Handbook of Industrial Organization*, 3 (eds. M. Armstrong and R. Porter), North-Holland.

[49] Tirole, J. (1988), The Theory of Industrial Organization, Cambridge, MA: MIT Press.

[50] 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.

[51] Villas-Boas, M. (2004), Communication Strategies and Product Line Design, *Marketing Science* 23(3), 304-316.

[52] Varian, H. (1980), A Model of Sales, *American Economic Review* 70(4), 651-659.

[53] Wernerfelt, B. (1986), Product Line Rivalry: Note, *American Economic Review* 76(4), 842-44.

[54] Wernerfelt, B. (1988), Umbrella Branding as a Signal of New Product Quality: An Example of Signalling by Posting a Bond, *RAND Journal of Economics* 19(3), 458-466.

[55] Wildenbeest, M. R. (2011), An Empirical Model of Search with Vertically Differentiated Products, *RAND Journal of Economics* 42(4), 729-757.

[56] Wolinsky, A. (1986), True Monopolistic Competition as a Result of Imperfect Information, *The Quarterly Journal of Economics* 101(3), 493-511.

[57] Zhou, J. (2014), Multiproduct Search and the Joint Search Effect, *American Economic Review*, 104(9), 2918-2939.