

DISCUSSION PAPER SERIES

DP15729
(v. 6)

Search and Price Discrimination Online

Eeva Muring

INDUSTRIAL ORGANIZATION

CEPR

Search and Price Discrimination Online

Eeva Muring

Discussion Paper DP15729
First Published 28 January 2021
This Revision 01 March 2022

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programmes:

- Industrial Organization

Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as an educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Eeva Muring

Search and Price Discrimination Online

Abstract

This paper theoretically studies price discrimination based on search costs. "Shoppers" have a zero and "nonshoppers" a positive search cost. A consumer faces a nondiscriminatory "common" price with some probability, or a discriminatory price. In equilibrium, firms mix over the common and the shoppers' discriminatory prices, but set a singleton nonshoppers' discriminatory price. Less likely price discrimination mostly benefits consumers. An individual firm's profit can increase in the number of firms. These results have important implications for regulations that limit the tracking of consumers (e.g., EU's GDPR, California's CCPA) and for evaluating competition online based on the number of firms.

JEL Classification: D83, D43

Keywords: cookies, consumer tracking, price discrimination, GDPR, Online markets, Imperfect Competition, sequential search

Eeva Muring - eevamuring@gmail.com
University of Bergen and CEPR

Search and Price Discrimination Online*

Eeva Mauring[†]

1st March 2022

Abstract

This paper theoretically studies price discrimination based on search costs. “Shoppers” have a zero and “nonshoppers” a positive search cost. A consumer faces a nondiscriminatory “common” price with some probability, or a discriminatory price. In equilibrium, firms mix over the common and the shoppers’ discriminatory prices, but set a singleton nonshoppers’ discriminatory price. Less likely price discrimination mostly benefits consumers. An individual firm’s profit can increase in the number of firms. These results have important implications for regulations that limit the tracking of consumers (e.g., EU’s GDPR, California’s CCPA) and for evaluating competition online based on the number of firms.

JEL codes: D83, D43

Keywords: *cookies, consumer tracking, price discrimination, GDPR, online markets, imperfect competition, sequential search*

1 Introduction

Increasingly more transactions are conducted online rather than offline. An important distinction between online and offline markets is that firms operating online can track consumers’ browsing behaviour more easily, for example, by using cookies. Firms use this consumer data to personalise their offers to consumers, in terms of content or prices.¹ Regulators both in the EU and U.S. seem to be concerned

*Acknowledgements: I thank Heski Bar-Isaac, Daniel García, Marc Goñi, Sander Heinsalu, Toomas Hinnosaar, Steffen Hoernig, Maarten Janssen, Bjørn Olav Johansen, Teis Lømo, Simon Martin, Clement Minaudier, Alexei Parakhonyak, Martin Peitz, Andrew Rhodes, Karl Schlag, Cole Williams, Jidong Zhou and audiences at Bergamo IO Workshop 2021, EARIE 2021, Paris Conference on Digital Economics, BEA Seminar, MaCCI/EPoS Virtual IO Seminar, Consumer Search Digital Seminar, DICE (Düsseldorf), OsloMet, the University of Leicester, NHH (Bergen), the University of Bergen, and the University of Vienna for many helpful questions and comments. Earlier title: “Sequential Search with Limited Price Discrimination”. Mauring acknowledges financial support from the Austrian Science Foundation FWF under project number FG 6-G.

[†]University of Bergen, University of Vienna and CEPR. Address: Department of Economics, University of Bergen, Fosswinckels gate 14, 5007 Bergen, Norway. eevamauring@gmail.com.

¹Many online firms’ privacy policies state that the firms use the consumer data they collect: to “personalize your experience” (Amazon, 2021), “create personalised Products” (Facebook, 2021), and “provide personalized services” (Google, 2021).

about these developments. The EU’s GDPR (General Data Protection Regulation) and California’s CCPA (California Consumer Privacy Act) make it easier for consumers to prohibit firms from accessing their data, for example, by simplifying the disabling of cookies.

This paper contributes to our understanding of online markets where firms use consumer data to price discriminate. It allows us to theoretically evaluate the effect of regulations that restrict firms’ possibilities to track consumers’ online behaviour and price discriminate based on this information. In particular, I study probabilistic price discrimination based on search costs in a sequential search model. My model predicts that regulations increase consumer welfare if they are strict enough. I also show that an increase in the number of firms may increase an *individual* firm’s profit, which erodes consumer welfare. This suggests that in online markets regulators should not interpret a large number of firms as a sign of healthy competition.

I model online markets as search markets because search costs online are considerable (De los Santos et al., 2012; Honka, 2014). Online firms can often infer something about the search costs of (potential) consumers. A person’s search cost can at least partly be inferred from his browsing behaviour: if he visits various online stores, spends lots of time there, and looks through scores of items, a firm that tracks him may reasonably infer that his search cost is low.² Inferring a person’s valuation, conversely, requires a firm to observe at least something about his past purchase history; information that can be obtained by fewer firms and at a higher cost. Empirical research has documented personalised prices based on consumers’ behaviour in several online markets, such as car and home insurance, hotels, flights, and car rental (FCA, 2019; Hannak et al., 2014; Ipsos et al., 2018).

In the model all consumers have a unit demand and the same valuation for a homogeneous good.³ They search sequentially for a low price. A fraction of the consumers (“shoppers”) have a zero search cost and the rest (“nonshoppers”) a positive search cost, as in Stahl (1989).⁴ Homogeneous firms compete over the consumers by setting prices.⁵ A firm can identify a consumer’s type (search cost) and offer him a discriminatory price, a “shoppers’ price” to a shopper and a

²Among other things, Amazon collects information on which products a user views or searches for, how long he spends on pages, and page interaction info like scrolling, clicks and even mouseovers (Amazon, 2021); p. 44 shows a screenshot. Both Facebook and Google collect data on a user’s interactions with its own and some third-party pages (Facebook, 2021; Google, 2021).

³I focus on the case of identical valuations to highlight the effects of price discrimination based on search cost only. I discuss heterogeneous valuations in Section 7.2.3.

⁴Stahl (1989) is the simplest model with heterogeneous search costs that captures the essence of sequential search in that some consumers’ outside option is determined endogenously.

⁵A price can more generally be thought of as a utility-transfer pair that a firm offers, where a different utility means a different quality or product match that the firm steers a consumer to.

“nonshoppers’ price” to a nonshopper, with some probability. This discrimination probability can depend on the consumer’s type. With the rest of the probability, the firm cannot identify the consumer’s type and offers him a nondiscriminatory “common price”. The discrimination probability can be seen as a proxy for the proportion of consumers who enable cookies or for firms’ average ability to track consumers, both of which decrease as privacy regulation becomes stricter.⁶

As the first main result, I derive and describe the model’s unique symmetric equilibrium. In principle, the mixed-strategy distributions of the three prices could be intricately linked, but turn out to be related quite simply. In equilibrium firms mix over common prices and shoppers’ prices, but set a singleton nonshoppers’ price. The dispersed common prices serve a dual purpose: attracting the shoppers on the one hand and extracting profits from the nonshoppers on the other. The highest shoppers’ price is equal to the lowest common price: shoppers identified as such always get a discount over the common prices and the size of the discount varies across firms. The highest common price and the nonshoppers’ price are equal to the nonshoppers’ cutoff price.

Second, I present novel comparative static results with respect to the price discrimination probabilities and the number of firms. In general consumers’ welfare is lower when price discrimination is more likely, but the details depend on whether the nonshoppers’ cutoff price is interior, i.e., below the valuation, or not. Consumer welfare and firms’ total profits are inversely related in the unique equilibrium because all consumers buy and nonshoppers search once each. If the nonshoppers’ cutoff price is interior, in general consumers suffer from more likely discrimination because firms can better target their price offers. Firms can, thus, raise all prices and extract more from consumers, especially nonshoppers. If, instead, the nonshoppers’ cutoff price is equal to the valuation, firms cannot raise the highest prices when price discrimination becomes more likely. The nonshoppers’ cutoff price is interior if the price discrimination probabilities are low enough. According to my model’s results, therefore, regulations that limit consumer tracking, such as the EU’s GDPR or California’s CCPA, certainly increase consumer welfare if they reduce the tracking of consumers enough or ease the disabling of cookies enough. In Section 6.1 I describe more precisely the effects of varying the price discrimination probability against one type of consumers at a time and argue that discrimination against shoppers drives the above result.

The most interesting comparative static result with respect to the number of

⁶The identification probabilities can be endogenised in several ways. I have chosen not to do so in the main model because different plausible extensions lead to different orderings of the identification probabilities; I discuss several possibilities in Sections 7.1 and 7.2.1.

firms is that an *individual* firm's profit may increase in the number of competitors. An increase in the number of firms leads to stiffer competition for both shoppers and nonshoppers. When the number of firms is large, however, an individual firm's chance of attracting many shoppers is small. So if the number of firms increases, a firm gives up on competing for shoppers and, instead, focuses on extracting rents from nonshoppers by raising some prices. These higher prices, in turn, relax competition and increase the firm's profit for some parameter values. This result is in contrast to Stahl (1989) where an individual firm's profit decreases in the number of firms. In my model, an individual firm benefits from more competitors if the extra competition works as a commitment device: induces the firm to refocus its efforts from competing fiercely for the price-sensitive consumers, the shoppers, to extracting more from the price-insensitive consumers, the nonshoppers. The extra flexibility lent by discriminatory prices allows the firm to do so in a manner that increases its profit. Since firms' profits can increase in the number of firms, my model suggests that regulators should not interpret a large number of sellers as a credible sign of healthy competition in online markets.

Another comparative static result with respect to the number of firms is that the dispersion of prices (as measured by the range of prices) increases in the number of firms. Many empirical papers document that price dispersion in online markets increases in the number of firms, while some find that the average price increases and others that it decreases in the number of firms (Haynes and Thompson, 2008; Grewal et al., 2010; Wang et al., 2021). A necessary condition for a firm's profit to increase in the number of firms in my model's equilibrium is that the average paid price increases in the number of firms. Thus, my model provides a possible explanation to the joint empirical findings that a higher number of firms online increases both price dispersion and the average price.

The model's results apply more generally, although I model the consumers in a very specific manner. Two features are important for the results. First, that consumers differ in price sensitivity. Second, that the prices that firms set to different consumer groups are related in equilibrium. These features generate the economic forces that underlie the comparative static results in my model. In particular, a firm's profit can increase in the number of competitors if the increase effectively functions as a commitment to not compete fiercely for the price-sensitive consumers. If all firms compete less for the price-sensitive consumers, the firms can focus on extracting profit from price-insensitive consumers by raising their prices. Since this slackens competition further, a firm can also raise the prices set to price-sensitive consumers and its profit increases.

Finally, I discuss several extensions to the model in Section 7. The first set

of extensions endogenise the consumer identification probabilities. The second set discusses more dimensions of consumer heterogeneity, including more general search cost distributions and dispersed valuations. Finally, I discuss correlated identification events. Some extensions may change the equilibrium characterisation somewhat: also nonshoppers' prices can be dispersed or a gap between the nonshoppers' price and the highest common price may arise. In all extensions, however, as long as some shoppers and nonshoppers are identified and others not, the same economic forces that generate my model's interesting results are at play.

Literature. My paper contributes to the literature on type-based price discrimination in imperfectly competitive markets. The closest papers, on consumer search, have focused on discrimination based on consumer valuations (Fabra and Reguant, 2020; Preuss, 2021) or third-degree price discrimination (Atayev, 2020). Bergemann et al. (2021) derive an upper bound on the prices at which trades occur for various information structures when firms can discriminate based on consumers' price count. In Armstrong and Vickers (2019) consumers exogenously know one price (captives) or multiple prices and firms discriminate against captives with probability zero or one. The only other paper with probabilistic price discrimination in an oligopoly setting that I know of is a duopoly model of Belleflamme et al. (2020). I focus on probabilistic price discrimination based on search costs and derive comparative statics that can be used to evaluate regulation.

My paper is also related to the literature on behaviour-based price discrimination, where a firm can set special prices to consumers who bought from it before.⁷ The more related papers are Chen (1997) and Taylor (2003), where a firm learns the switching cost of a consumer. In the search model in Armstrong and Zhou (2016), firms set different prices to new and returning consumers. Conceptually, discriminating between consumers based on their short purchase histories is akin to discriminating based on coarse signals about consumer types, and discriminating based on long histories akin to precise signals about consumer types. My model belongs to the latter group because I aim to model the idea that storing long histories has become technologically feasible and cheaper.

Finally, my paper is related to papers on online privacy.⁸ The most related is Braghieri (2019) where firms can discriminate between consumers based on whether they search or not. Braghieri (2019) focuses on the consumers' privacy choices rather than assessing the effects of regulation.

⁷See, for example, Hart and Tirole (1988) for a monopoly model or Villas-Boas (1999) and Fudenberg and Tirole (2000) for imperfect competition. In some papers in this literature, firms discriminate also based on the consumer's type.

⁸See, for example, Acquisti and Varian (2005), Conitzer et al. (2012), Montes et al. (2019), Ichihashi (2020), Hidir and Vellodi (2021), Loertscher and Marx (2020), and Braghieri (2019).

Section 2 introduces the model, and Sections 3 and 4 the problems of consumers and firms respectively. Section 5 describes the equilibrium and Section 6 the comparative statics' results. Section 7 discusses several extensions and Section 8 concludes. All proofs are in the Appendix.

2 Model

The model is a unit-demand version of Stahl (1989) with probabilistic price discrimination.

Consumers. A measure one of consumers look for a homogeneous good at a low price to maximise utility. Each consumer has a unit demand and values the good at v .⁹ Consumers are one of two types: a fraction $\lambda > 0$ of the consumers are “shoppers” and a fraction $1 - \lambda$ are “nonshoppers”. Shoppers have a zero search cost and see the price that they would be charged at each firm before deciding which firm to buy from. Nonshoppers do not know the prices ex ante and search through the firms in a random order. When a nonshopper visits firm i , he finds out what price firm i offers him. He pays a small search cost of $\alpha > 0$ to find out each price offer (except for the first).¹⁰ Recall is free as usual. A consumer can exit the market if his expected value from participating is negative. Total consumer welfare is measured by the sum of the net utilities across the buyers who buy minus the total search costs paid by nonshoppers.

Firms. Each of the $N \geq 2$ firms produces the homogeneous good at zero marginal cost. Firms set prices and maximise expected profits. A price can more generally be thought of as a utility-transfer pair that a firm offers, where the utility that a firm steers a consumer to (derived from, e.g., quality or product match) is produced at a linear marginal cost. Firms can price discriminate with some probability in $[0, 1)$. In particular, if a shopper (nonshopper) visits firm i , then i identifies the consumer's type with probability μ_s (μ_n); μ_s and μ_n may differ. If firm i identifies the consumer's type, it can offer the consumer a discriminatory price: its “shoppers' price” p_s^i to a shopper and its “nonshoppers' price” p_n^i to a nonshopper. If the firm cannot identify the consumer's type, it must offer him its “common price” p_c^i . The firm's price offers to all consumers are independent: one shopper may be offered p_s^i at firm i , whereas another shopper or a nonshopper may be offered p_c^i at i .¹¹ The price offers are also independent across firms: a

⁹I discuss heterogeneous valuations in Section 7.

¹⁰First search is free in Stahl (1989) and in most papers that build on it.

¹¹The assumption that different shoppers (nonshoppers) get independent price offers at a firm is without loss of generality. The equilibrium that I derive is the unique symmetric equilibrium also if all consumers with the same search cost are offered the same price at a firm.

shopper may be offered p_c^i at firm i and p_s^j at firm j .¹² A firm does not see which prices are offered by other firms (but knows their equilibrium strategies as usual). If $\mu_s = \mu_n = 0$, my model collapses to the unit-demand version of Stahl (1989).

The setup can be interpreted as follows. With probability μ_s a shopper enables firm i 's cookies that track his behaviour online (and is offered a personalised price p_s^i) and with the rest of the probability he disables firm i 's cookies (and is offered a nonpersonalised price p_c^i). A nonshopper enables cookies of a firm with a potentially different probability μ_n than a shopper.¹³ In reality, a firm can access information about a person's interactions with its own website (immediately or at the person's later visit) using first-party cookies and with other firms' websites by using third-party cookies.¹⁴

Timing. First, a firm sets its prices p_c , p_n and p_s . Second, each shopper gets price offers at all firms and decides which firm to buy from. Third, each nonshopper randomly chooses a firm to visit first and learns his price offer at the firm. He decides whether to accept the price offer and buy or to continue costly search. If he continues, he draws the next firm to visit at random from amongst the previously unvisited firms. Fourth, utilities are realised.

Strategies. A firm's pure strategy is a triple of real numbers (p_c, p_s, p_n) : a common price p_c , a shoppers' price p_s , and a nonshoppers' price p_n . A firm's mixed strategy, F , is a joint probability distribution over its pure strategies. A nonshoppers' strategy specifies at which prices to buy and at which to continue searching. A shopper's strategy specifies at which firm to buy. I assume that a consumer accepts a current price offer when just indifferent.

Equilibrium. I focus on symmetric equilibria. In equilibrium, an agent plays an optimal strategy, taking as given all other agents' behaviour. Anticipating that a firm optimally chooses independent p_c , p_s , and p_n , I denote the marginal mixed-strategy distribution of p_j by $F_j(p_j)$, expectations with respect to F_j by \mathbb{E}_j , and the lowest and highest prices in the support of F_j by \underline{p}_j and \bar{p}_j respectively for $j = c, s, n$. In a symmetric equilibrium, all firms use the same mixed-strategy distributions F_c , F_s , and F_n . In a symmetric equilibrium all shoppers and nonshoppers use the same optimal policies. As is standard in sequential search models, I assume that nonshoppers hold passive beliefs: if they observe a deviation by one firm, they believe that no other firm has deviated.

¹²I discuss correlated identification events in Section 7.3.

¹³I discuss privacy choices in Section 7.1.1.

¹⁴For example, when a person visits Amazon's website, Amazon uses first-party cookies, but also allows other firms to place (third-party) cookies on the person's browser (Amazon, 2021); p. 44 shows a screenshot.

3 Consumers' problems

I briefly describe the consumers' problems. A nonshopper's problem is almost standard. Suppose that the nonshopper is visiting the k th firm, $k < N$, and that the lowest price that he has seen so far is p . He does not per se care about whether p is a firm's common or nonshoppers' price. Should the nonshopper accept the price p or continue to search? A sufficient condition for continuing to search to be optimal is that the cost of searching one more firm, α , is less than its benefit, $B(p)$, because the nonshopper can always exit after searching the next firm.

If nonshoppers expect that each firm chooses prices according to distribution F , the expected benefit of searching one more firm is

$$B(p) = \mu_n \mathbb{E}_F[p - p_n | p_n < p] + (1 - \mu_n) \mathbb{E}_F[p - p_c | p_c < p].$$

With probability μ_n , the next firm identifies the nonshopper and offers him its nonshoppers' price p_n so his benefit is $p - p_n$ if $p_n < p$. With probability $1 - \mu_n$, the next firm does not identify the nonshopper so his benefit is $p - p_c$ if $p_c < p$. A sufficient condition for continuing to search to be optimal is, thus, that $B(p) > \alpha$.

Two observations complete the description of a nonshopper's optimal policy. First, the benefit of continuing to search $B(p)$ increases in p and $B(\min\{\underline{p}_n, \underline{p}_c\}) = 0$ so the equation $B(p) = \alpha$ has a unique solution. Second, a consumer can search more than one additional firm only if he searches the next one. Thus, a necessary condition for searching one or more additional firms to be optimal is that the cost of searching one more firm is less than its benefit.

Altogether, the nonshoppers' optimal search rule is a cutoff policy: to accept the first price offer that falls below the optimal cutoff price, ϕ_n , and to continue searching otherwise. The optimal cutoff ϕ_n solves

$$B(\phi_n) = \alpha, \tag{1}$$

if the solution is below the valuation v . If the solution to equation (1) exceeds v , the optimal cutoff is $\phi_n = v$. If the nonshopper gets price offers that exceed his cutoff price at all firms, then after visiting the last firm he accepts the lowest price he saw if it is below v . In a symmetric equilibrium, all nonshoppers use the same cutoff price. I derive the cutoff price ϕ_n explicitly in the proof of Proposition 1.

A shopper's problem is standard. A shopper observes all firms' price offers before deciding which firm to buy from so he buys at the firm that offers him the lowest price as long as this is below v . If all offers exceed v , he exits the market.

4 A firm's problem

I set up a firm's problem and show that it optimally sets independent prices. Suppose that firm i has decided to set prices (p_c^i, p_s^i, p_n^i) as its common, shoppers' and nonshoppers' prices respectively. Firm i maximises its expected profit. Instead of considering the firm's ex ante expected profit, let us first think about the firm's interim profit: the expected profit when the firm knows that the visiting consumer is a nonshopper, a shopper, or an unidentified consumer.

Suppose first that firm i identifies a visiting consumer as a nonshopper, thus, offers him its nonshoppers' price p_n^i . Firm i 's interim profit from the nonshopper is

$$\pi_n^i(p_n^i) = \begin{cases} p_n^i & \text{if } p_n^i \leq \phi_n, \\ p_n^i P(n \text{ returns to } i | p_n^i) & \text{if } p_n^i \in (\phi_n, v]. \end{cases} \quad (2)$$

The first line of equation (2) says that the identified nonshopper buys immediately from firm i if p_n^i is weakly below his cutoff price. The second line says that the nonshopper buys at a price that exceeds his cutoff price only if he returns to firm i after visiting all firms and encountering prices there that exceed ϕ_n , with probability $P(n \text{ returns to } i | p_n^i)$. This probability depends on p_n^i and the prices that the nonshopper is offered at firms other than i , but does not depend on firm i 's common and shoppers' prices p_c^i and p_s^i because the identified nonshopper cannot get a different price offer at i than p_n^i .

Suppose now that firm i identifies a visiting consumer as a shopper, thus, offers him its shoppers' price p_s^i . Firm i 's interim profit from the shopper is

$$\pi_s^i(p_s^i) = D_s^i(p_s^i) p_s^i, \quad (3)$$

where $D_s^i(p_s^i)$, which I address shortly, is the probability that a shopper who is offered price p_s^i at firm i buys at i .

Finally, suppose that firm i cannot identify the type of the visiting consumer, thus, offers him its common price p_c^i . Firm i 's interim profit from such a consumer is

$$\pi_c^i(p_c^i) = \begin{cases} [1 - \nu(s|\text{unid})] p_c^i + \nu(s|\text{unid}) D_s^i(p_c^i) p_c^i & \text{if } p_c^i \leq \phi_n, \\ [1 - \nu(s|\text{unid})] p_c^i P(n \text{ returns to } i | p_c^i) + \nu(s|\text{unid}) D_s^i(p_c^i) p_c^i & \text{if } p_c^i > \phi_n, \end{cases} \quad (4)$$

where $\nu(s|\text{unid})$, that I derive below, denotes firm i 's interim belief that the unidentified consumer is a shopper given that the consumer visits i for the first time and that i cannot identify his type. The first terms on both lines of equation (4) are the

expected profits from a nonshopper: if the unidentified consumer is a nonshopper, he buys immediately from firm i if p_c^i does not exceed his cutoff price ϕ_n and buys from firm i at price $p_c^i > \phi_n$ only after visiting all firms and encountering prices there that exceed ϕ_n , with probability $P(n \text{ returns to } i | p_c^i)$. The second terms in equation (4) are the expected profits from a shopper: if the unidentified consumer is a shopper, he buys from firm i at price p_c^i with probability $D_s^i(p_c^i)$. As above, $P(n \text{ returns to } i | p_c^i)$ depends on p_c^i and the prices that the nonshopper is offered at firms other than i , but does not depend on firm i 's other prices p_n^i and p_s^i .

A shopper who is offered price p^i at firm i buys at i with probability $D_s^i(p^i)$ in equations (3) and (4). The probability depends on firm i 's price distribution via p^i and on the other firms' price distributions because firm i sells to a visiting shopper at p^i only if p^i is lower than the shopper's offers at all other firms. However, the other prices of firm i , p_s^i and p_n^i if $p^i = p_c^i$ or p_c^i and p_n^i if $p^i = p_s^i$, do not affect $D_s^i(p^i)$ because a shopper only gets one price offer from each firm. Firm i does not know exactly if any firm $j \neq i$ offers the given shopper its shoppers' or common price, but i knows that j offers p_s^j if j can identify the shopper's type (with probability μ_s) and p_c^j otherwise. Also, i knows that the price offers at any two firms are independent. Thus, knowing the price distributions of other firms, firm i can calculate the expected probability of selling to a shopper at price p^i , $D_s^i(p^i)$. The probabilities $P(n \text{ returns to } i | p^i)$ and $D_s^i(p^i)$ do not have nice closed-form solutions at this level of generality. I derive the probabilities explicitly for the unique symmetric equilibrium in the proof of Proposition 1.

Firm i 's interim belief about the type of a visiting unidentified consumer can be derived using Bayes' rule. In terms of the odds ratio, we get

$$\frac{\nu(s|\text{unid})}{1 - \nu(s|\text{unid})} = \frac{P(s)}{P(n)} \frac{P(\text{visits } i | s)}{P(\text{visits } i | n)} \frac{P(\text{unid} | s, \text{visits } i)}{P(\text{unid} | n, \text{visits } i)},$$

where s (n) stands for the event that a consumer is a shopper (nonshopper) and "unid" for the event that the consumer is unidentified. The prior probability that a visiting consumer is a shopper is λ and the probability that he is not identified is $1 - \mu_s$. The prior probability that a visiting consumer is a nonshopper is $1 - \lambda$ and the probability that he is not identified is $1 - \mu_n$. We, thus, only need to figure out the probabilities that a consumer of a given type visits firm i for the first time (i.e., excluding returns by nonshoppers after searching through all firms), $P(\text{visits } i | s)$ and $P(\text{visits } i | n)$.

Shoppers visit all firms before deciding where to buy so $P(\text{visits } i | s) = 1$. Nonshoppers search randomly so the probability that a nonshopper visits firm i first is $\frac{1}{N}$. If firms $j \neq i$ offer only prices weakly below the nonshoppers' cutoff

price ϕ_n , no nonshopper that visited firm $j \neq i$ first visits also firm i so the total probability that a nonshopper visits firm i for the first time is $P(\text{visits } i|n) = \frac{1}{N}$. If, instead, some firms offer with positive probability prices that exceed ϕ_n , firm i can also be visited for the first time by a nonshopper as the k th firm in the nonshopper's sequence of visits for $k = 2, \dots, N$. However, note that the total probability

$$P(\text{visits } i|n) = \sum_{k=1}^N P(\text{visits } i \text{ as } k\text{'th firm}|n),$$

does not depend on prices at firm i (but does depend on other firms' prices) because nonshoppers' search is random. Thus, firm i 's posterior belief about the type of a visiting unidentified consumer is

$$\frac{\nu(s|\text{unid})}{1 - \nu(s|\text{unid})} = \frac{\lambda}{1 - \lambda} \frac{1}{P(\text{visits } i|n)} \frac{1 - \mu_s}{1 - \mu_n},$$

which is independent of (p_c^i, p_s^i, p_n^i) .

In the interim profits of firm i , equations (2), (3) and (4), its prices p_n^i , p_s^i , and p_c^i appear in turn, but never simultaneously. The only additional terms that appear in firm i 's ex ante expected profit are the expected number of consumers of each type that visit i for the first time. These are independent of firm i 's prices because i 's prices affect neither the total number of consumers of a type, the probability with which i identifies a consumer nor, as I argued above, the probabilities that a consumer of a given type visits firm i for the first time. Thus, considering firm i 's interim profits shows that the firm's optimal prices are set independently of each other: in any equilibrium F_c , F_s and F_n are independent.

5 Equilibrium

In this Section, I describe the equilibrium price distributions and their novel features. The equilibrium is summarised in

Proposition 1. *In the unique symmetric equilibrium,*

- *a firm's strategy comprises*

(i) *the distribution of common prices $F_c(p_c) = 1 - \left(\frac{\bar{p}_c - p_c}{\gamma N p_c}\right)^{\frac{1}{N-1}}$, with support $[\underline{p}_c, \bar{p}_c]$ where $\bar{p}_c = \min \left\{ \frac{\alpha}{1 - \mu_n} \left[1 - \int_0^1 (1 + \gamma N y^{N-1})^{-1} dy \right]^{-1}, v \right\}$, $\underline{p}_c = \frac{\bar{p}_c}{1 + \gamma N}$, and $\gamma := \frac{\lambda(1 - \mu_s)^N}{(1 - \lambda)(1 - \mu_n)}$,*




Figure 1: The equilibrium distribution of shoppers' prices f_s (red dashed) and of common prices f_c (black solid) in my model for $\mu_s = \mu_n = \frac{1}{4}$, and the equilibrium distribution of prices in Stahl (1989) (grey dotted); $N = 5$, $\lambda = \frac{1}{2}$, $\alpha = \frac{1}{20}$, $v = 1$.

(ii) the distribution of shoppers' prices $F_s(p_s) = \frac{1}{\mu_s} \left[1 - (1 - \mu_s) \left(\frac{\bar{p}_s}{p_s} \right)^{\frac{1}{N-1}} \right]$,
with support $[\underline{p}_s, \bar{p}_s]$ where $\bar{p}_s = \underline{p}_c$ and $\underline{p}_s = (1 - \mu_s)^{N-1} \bar{p}_s$, and

(iii) the nonshoppers' price $p_n = \phi_n$.

- a nonshoppers' strategy is to accept all prices as $p_n = \bar{p}_c = \phi_n$.
- a shoppers' strategy is to buy at the firm that offers him the lowest price.

In the proof of Proposition 1 I first derive some necessary conditions on the equilibrium price distributions, then show that the behaviour described in the Proposition is an equilibrium, and last show that it is the unique symmetric equilibrium. In the proof, I have to take into account, first, that shoppers see some common and some shoppers' prices. Second, that the supports of the common and shoppers' prices could overlap partly, fully, or not at all. I show that the two distributions' supports just touch.

Figure 1 plots an example of the equilibrium pdfs of shoppers' prices f_s (red dashed) and of common prices f_c (black solid) for $\mu_s = \mu_n = \frac{1}{4}$, $N = 5$, $\lambda = \frac{1}{2}$, $\alpha = \frac{1}{20}$, and $v = 1$. For these parameter values, the nonshoppers' cutoff price is interior: $\phi_n = 0.37$. For comparison, Figure 1 also plots the standard single-unit Stahl (1989) solution, i.e., f_c for $\mu_s = \mu_n = 0$ (grey dotted).

I now describe several aspects of the equilibrium. The two extreme cases of no price discrimination ($\mu_s = \mu_n = 0$) and full price discrimination ($\mu_s = \mu_n = 1$) are well-known benchmarks. Without price discrimination, we have Stahl

(1989) and only common prices are dispersed (because shoppers' prices are never offered). With full price discrimination, common prices collapse to the valuation (because close to $\mu_s = 1$, common prices are used primarily to extract profits from nonshoppers), while the shoppers' prices concentrate at the Bertrand outcome, price zero. Both the common and shoppers' prices are dispersed for all interior price discrimination probabilities (in fact, $\mu_s \in (0, 1)$ suffices).

Common prices in my model serve a dual purpose: attracting shoppers on the one hand and extracting surplus on the other. The expression for the common price distribution F_c is similar to the price distribution in the unit-demand version of Stahl (1989), except that γ replaces $\frac{\lambda}{1-\lambda}$. In my model γ captures the relative importance of shoppers and nonshoppers among the consumers for whom a firm's common price is relevant: the nonshoppers who are unidentified by the firm (i.e., a fraction $1 - \mu_n$) and the shoppers who are unidentified by all firms (i.e., a fraction $(1 - \mu_s)^N$). As an aside, note that μ_s and μ_n enter γ differently: the probabilities of price discrimination against shoppers and nonshoppers affect the equilibrium distributions differently, as I describe in detail in Section 6.

The shoppers' price distribution, the most interesting new part of the equilibrium, has the following features. Unlike the pdf of common prices, the pdf of shoppers' prices decreases. If a firm sets a low rather than a high shoppers' price, it attracts some more shoppers, but gets less revenue on all served shoppers. The firm is indifferent between a low and high shoppers' price if it sets the low price more frequently. Firms optimally separate the common and shoppers' prices so that all shoppers' prices are below the common prices: $\bar{p}_s = \underline{p}_c$.¹⁵ The intuition is as follows. First, it is never optimal for a firm to set its common price lower than its shoppers' price. A low common price is wasteful if offered to a nonshopper and rarely attracts a shopper so the firm is better off by setting its shoppers' price below its common price instead. Second, there is no gap between the supports of the shoppers' and common price distributions. If there was a gap, a firm that is supposed to offer the highest shoppers' price would rather offer a shoppers' price between \bar{p}_s and \underline{p}_c : its demand would not be affected, but the revenue from identified shoppers would increase.

Since the shopper's price is below the common price at any firm, a consumer who is identified as very price-sensitive from cookies gets a discount over a consumer who disables cookies. An identified nonshopper, instead, gets a price offer above the common price at a firm. In other words, consumers who are identified as price-insensitive from cookies are always asked higher prices than consumers who disable cookies. Both a low shoppers' price and a high nonshoppers' price can

¹⁵See Lemma 6, p. 33, for the formal proof.

	Interior ϕ_n				Boundary ϕ_n		
	\bar{p}_c	\underline{p}_c	\underline{p}_s	π	\underline{p}_c	\underline{p}_s	π
$\mu \uparrow$	\uparrow	\uparrow	\uparrow	\uparrow	\uparrow	\downarrow	(\uparrow)
$\mu_s \uparrow$	\uparrow	\uparrow	\uparrow	\uparrow	\uparrow	(\uparrow)	(\uparrow)
$\mu_n \uparrow$	\uparrow	\downarrow	\downarrow	(\uparrow)	\downarrow	\downarrow	\downarrow
$N \uparrow$				(\uparrow)			\downarrow

Table 1: *New comparative static results. The joint price discrimination probability μ satisfies $\mu_s = \mu_n = \mu$. An increase for all relevant parameter values is denoted by \uparrow and for a nonempty open subset of parameter values by (\uparrow) .*

be implemented in a hidden manner, e.g., by a pop-up-window offering a discount to a shopper and by a higher delivery fee charged to a nonshopper.

6 Comparative statics

I present the comparative statics with respect to the price discrimination probabilities, which are all new, and highlight a new one with respect to the number of firms.¹⁶ The comparative statics with respect to the search cost and the fraction of shoppers are as in Stahl (1989).

The comparative statics depend on whether the nonshoppers' cutoff price ϕ_n is below or equal to their valuation. The solution to ϕ_n is interior if the price discrimination probabilities and the number of firms are low enough: if $\mu_s < \bar{\mu}_s$, $\mu_s < \bar{\mu}_n$ (or $\mu < \bar{\mu}$), and $N < \bar{N}$.¹⁷ Each of these critical parameter values depends on the values of the other parameters, but I suppress this relation in the notation for brevity. A selection of the comparative statics are summarised in Table 1, where the most interesting results, with respect to a firm's profit, are in boldface. The results are formalised and explained in turn in the next subsections.

6.1 Likelier price discrimination

Here I analyse the effect of varying the price discrimination probabilities. A decrease in these probabilities can be seen as the introduction or strengthening of consumer privacy protection regulations, such as the EU's GDPR and California's CCPA, because the regulations make it harder for firms to track consumers.

¹⁶More comparative static results with respect to the number of firms are in Proposition 4.

¹⁷The solution is interior if the parameters are low enough because an interior ϕ_n increases in $x = \mu_s, \mu_n, \mu, N$; see Propositions 3 and 4. I define \bar{x} as the smallest x such that $\phi_n = v$.

I first analyse the effect of varying the discrimination probabilities against shoppers and nonshoppers simultaneously. To do so, I set the discrimination probabilities against all consumers to be the same $\mu := \mu_s = \mu_n$ and call μ the joint price discrimination probability. The comparative static result with respect to μ is the relevant one if we want to study the effect of introducing or tightening a consumer privacy regulation and do not have a compelling reason to think that shoppers and nonshoppers are discriminated against with different frequencies. I discuss why the frequencies could differ on p. 19 and Sections 7.1 and 7.2.1.

The comparative statics with respect to the joint price discrimination probability are formalised in

Proposition 2. *Let the probabilities of price discrimination against shoppers and nonshoppers be the same $\mu = \mu_s = \mu_n$. If the joint probability of price discrimination, μ , increases,*

- (i) *and $\mu < \bar{\mu}$ (i.e., ϕ_n has an interior solution), then the lowest and highest common prices, \underline{p}_c and \bar{p}_c , the lowest shoppers' price, \underline{p}_s , and the expected profit, π , increase.*
- (ii) *and $\mu \geq \bar{\mu}$ (i.e., ϕ_n has a boundary solution), then the lowest common price, \underline{p}_c , increases; the lowest shoppers' price, \underline{p}_s , decreases; and the expected profit, π , increases for $\mu < \check{\mu}$ and decreases for $\mu \geq \check{\mu}$.¹⁸*

In sum, Proposition 2 states that more likely price discrimination mostly raises prices, benefits firms, and hurts consumers. Consumers are hurt if firms' total profits increase: in equilibrium all consumers buy and the nonshoppers search exactly once so total consumer welfare is measured by the negative of the firms' total profits. Figure 2 illustrates how the joint probability of price discrimination affects the bounds of the price distributions.

More specifically, Proposition 2 states that if the nonshoppers' cutoff price is interior ($\phi_n < v$) and price discrimination becomes likelier, the lowest and highest prices increase. This, not surprisingly, increases a firm's profit and lowers the consumer welfare. Intuitively, if μ increases, then firms quote the discriminatory prices more often to both types of consumers. At the same time, firms can use their common prices to target one type of consumers more, and they raise the common prices to use them more for extracting profits from the price-insensitive consumers, the nonshoppers. This relaxes competition, all prices rise, profit increases, and consumer welfare falls.

¹⁸The critical values $\bar{\mu}$ and $\check{\mu}$ are not related; $\bar{\mu}$ depends on v and α while $\check{\mu}$ does not.




Figure 2: *The bounds of the equilibrium distributions, \bar{p}_c (purple solid), $\bar{p}_s = \underline{p}_c$ (blue dashed) and \underline{p}_s (green dotted), in the joint probability of price discrimination $\mu = \mu_s = \mu_n$; $N = 5$, $\lambda = \frac{1}{2}$, $\alpha = \frac{1}{20}$, $v = 1$.*

If, instead, the nonshoppers' cutoff price has a boundary solution ($\phi_n = v$) and price discrimination becomes likelier, the lowest shoppers' prices decrease. Now firms cannot raise the highest common price (because $\bar{p}_c = \phi_n = v$), but they still raise the lowest common price to extract more from nonshoppers. But because they also compete more often for shoppers via the shoppers' prices, they lower the lowest shoppers' prices. In total, the higher lowest common prices do not always make up for the lost profit from the lower shoppers' prices and a firm's profit can decrease as price discrimination becomes likelier.

Note that a firm's profit when price discrimination is impossible is lower than when price discrimination is perfect: $\pi(\mu = 0) < \pi(\mu = 1)$.¹⁹ Together with part (i) of Proposition 2, this implies that there exists a price discrimination probability $\hat{\mu} \in (0, \bar{\mu})$ such that for all $\mu < \hat{\mu}$, we have $\pi(\mu) < \pi(\hat{\mu})$, and vice versa for all $\mu > \hat{\mu}$. In other words, consumer welfare is higher at all discrimination probabilities below $\hat{\mu}$ than at all discrimination probabilities above $\hat{\mu}$.

Do regulations that give consumers more rights over their data, such as the EU's GDPR or California's CCPA, raise consumer welfare? According to my model, the answer depends partly on how often firms price discriminated prior to the regulation's enforcement. To be sure that a regulation raises consumer welfare, it must reduce the tracking of consumers enough or ease the disabling of cookies enough (i.e., push the discrimination probability to $\hat{\mu}$ or lower).

The underlying forces behind these results become clearer if we allow the price

¹⁹The expected equilibrium profit of a firm is $\pi = \frac{1-\lambda}{N} \bar{p}_c + \mu \lambda \underline{p}_s$.

discrimination probabilities against shoppers and nonshoppers to differ. The comparative statics with respect to these probabilities are formalised in

Proposition 3. *If the probability of price discrimination against shoppers, μ_s , increases*

(i) *and $\mu_s < \bar{\mu}_s$ (i.e., ϕ_n has an interior solution), then the lowest and highest common prices, \underline{p}_c and \bar{p}_c , the lowest shoppers' price, \underline{p}_s , and the expected profit, π , increase.*

(ii) *and $\mu_s \geq \bar{\mu}_s$ (i.e., ϕ_n has a boundary solution), then the lowest common price, \underline{p}_c , increases; the lowest shoppers' price, \underline{p}_s , increases for $\mu_s < \check{\mu}_{s1}$ and decreases for $\mu_s \geq \check{\mu}_{s1}$; and the expected profit, π , increases for $\mu_s < \check{\mu}_{s2}$ and decreases for $\mu_s \geq \check{\mu}_{s2}$ with $\check{\mu}_{s1} < \check{\mu}_{s2}$.*

If the probability of price discrimination against nonshoppers, μ_n , increases

(i) *and $\mu_n < \bar{\mu}_n$ (i.e., ϕ_n has an interior solution), then the highest common price, \bar{p}_c , increases, but the lowest common price, \underline{p}_c , and the lowest shoppers' price, \underline{p}_s , decrease. A sufficient condition for the expected profit, π , to increase is that $\mu_n < \check{\mu}_{n1}$.*

(ii) *and $\mu_n \geq \bar{\mu}_n$ (i.e., ϕ_n has a boundary solution), then the lowest common price, \underline{p}_c , the lowest shoppers' price, \underline{p}_s , and the expected profit, π , decrease.*

In sum, Proposition 3 shows that the probabilities of price discrimination against shoppers and nonshoppers have different effects on the market outcomes. Roughly speaking, more likely price discrimination only against shoppers mostly raises prices and benefits firms, whereas more likely discrimination only against nonshoppers lowers some prices and benefits firms less frequently. Figures 3a and 3b illustrate how the probability of price discrimination against shoppers and nonshoppers respectively affects the bounds of the price distributions. Comparing them to Figure 2 suggests that the effects of the joint price discrimination probability are driven by discrimination against shoppers rather than nonshoppers. Discrimination against shoppers has a larger effect because firms compete fiercely for shoppers, but have temporary monopoly power over nonshoppers.

What drives the different effects? Consider an increase in the probability of price discrimination against shoppers, μ_s . A firm now offers, on the one hand, its common price more often to nonshoppers and, on the other, to a shopper more often its shoppers' price. These two effects are quite different. The first effect means that the firm wants to quote a high common price more frequently in order

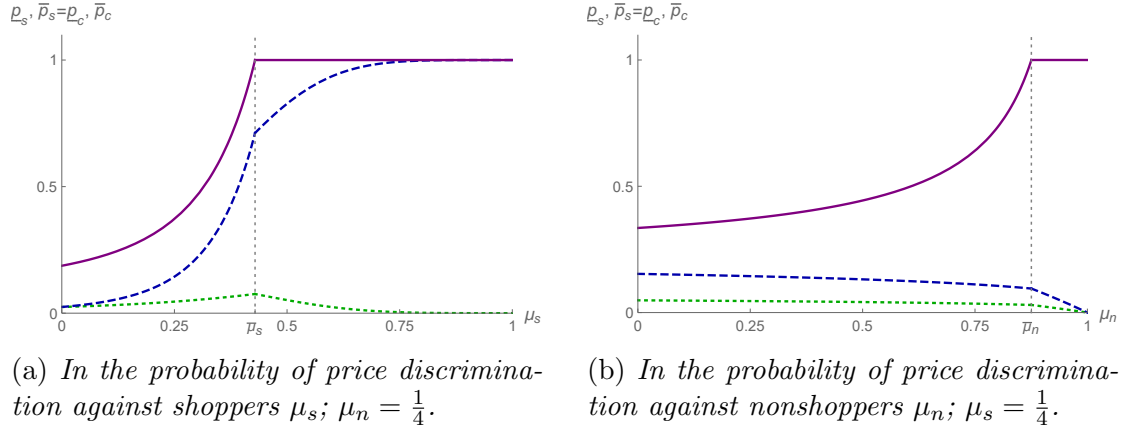


Figure 3: The bounds of the equilibrium distributions, \bar{p}_c (purple solid), $\bar{p}_s = \underline{p}_c$ (blue dashed) and \underline{p}_s (green dotted), in the probabilities of price discrimination; $N = 5$, $\lambda = \frac{1}{2}$, $\alpha = \frac{1}{20}$, $v = 1$.

to extract more from nonshoppers; this pushes towards higher common prices and the weakening of competition. The second effect means that the firm competes for the shoppers more often via its low shoppers' price; this pushes towards lower shoppers' prices and the stiffening of competition. If firms are not restricted in raising the highest common price, the first effect always dominates: all prices and a firm's profit increase in μ_s . But if the nonshoppers' cutoff price, thus, the highest common price, is equal to the valuation, the second effect can dominate: both the lowest shoppers' price and profit can decrease in μ_s .

Now consider an increase in the probability of price discrimination against nonshoppers, μ_n . Again, this has two effects: a firm now offers, on the one hand, its common price more often to shoppers and, on the other, to a nonshopper more often its nonshoppers' price. The first effect means that the firm competes for the shoppers more often via the lowest common prices; this pushes towards lower shoppers' prices and the stiffening of competition. But the second effect means that the firm wants to raise its nonshoppers' price, or, equivalently, the highest common price, to extract more from nonshoppers; this pushes towards higher highest common prices and the weakening of competition. If the firm is not restricted in raising the highest common price, the second effect is mostly strong enough so that competition softens and a firm's profit increases in μ_n . But if the nonshoppers' cutoff price has a boundary solution, the first effect dominates: both the lowest prices and profit decrease in μ_n .

All in all, if a regulation lowers the price discrimination probabilities simultaneously, then its effect depends partly on how often firms discriminated prior to the regulation's enforcement and partly on which price discrimination probability is larger. If $\mu_s > \mu_n$, the nonshoppers' cutoff price is more likely to have

a boundary solution where firms can benefit if price discrimination becomes less likely. A shopper is more likely to be identified than a nonshopper, for example, if a shopper leaves his virtual footprints in more online places than a nonshopper by visiting many websites (Section 7.1.2 discusses this more). If, conversely, $\mu_s < \mu_n$, the nonshoppers' cutoff price is more likely to have an interior solution where firms suffer and consumers benefit if price discrimination becomes less likely. A shopper is less likely to be identified than a nonshopper, for example, if shoppers are more tech-savvy or computer-literate than nonshoppers: shoppers are better both at online shopping and at disabling cookies (Section 7.2.1 discusses this more).

6.2 More firms

The comparative statics with respect to the number of firms are formalised in

Proposition 4. *If the number of firms, N , increases,*

- (i) *and $N < \bar{N}$ (i.e., ϕ_n has an interior solution), then the highest common price, \bar{p}_c , increases. A sufficient condition for the lowest common price, \underline{p}_c , to increase is that $N > \check{N}_1$. The lowest shoppers' price, \underline{p}_s , increases and decreases for nonempty open sets of parameter values. Sufficient conditions for the expected profit, π , to increase are $N > 4$, $\mu_s \geq \check{\mu}_{s3}$, and $\mu_n \in (\check{\mu}_{n2}, \check{\mu}_{n3})$. The dispersion of prices, $\bar{p}_c - \underline{p}_s$, increases.*
- (ii) *and $N \geq \bar{N}$ (i.e., ϕ_n has a boundary solution), then the lowest common price, \underline{p}_c , increases for all $N > \check{N}_1$ and decreases for all $N \leq \check{N}_1$; the lowest shoppers' price, \underline{p}_s , and the expected profit, π , decrease; the dispersion of prices, $\bar{p}_c - \underline{p}_s$, increases.*

The bounds of the price distributions change in the number of firms similarly to those in Stahl (1989). If the nonshoppers' cutoff price is interior and the number of firms increases, the highest common price increases, but for some parameter values the lowest shoppers' price decreases. If the number of firms increases, then competition becomes fiercer. In response, a firm sometimes offers lower lowest prices to the shoppers. But a firm would have to lower the shoppers' prices a lot in order to attract many shoppers. It is, instead, more profitable to focus on extracting profit from nonshoppers and increase the higher common prices. The increase in the highest common price is so large that the dispersion of prices increases in the number of firms.

In Stahl (1989) a firm's profit always decreases in the number of firms because the positive effect of an increase in prices is outweighed by the increase in




Figure 4: An individual firm's expected profit π in the number of firms N , in my model for $\mu_s = \mu_n = \frac{1}{4}$ (red solid) and in Stahl (1989) (grey dashed); $\lambda = \frac{1}{2}$, $\alpha = \frac{1}{20}$, $v = 1$.

competition.²⁰ In my model, in contrast, an *individual firm's profit* increases in the number of firms for a nonempty open set of parameter values. The reason is the following. If there are relatively many firms, then they compete for shoppers almost solely via the shoppers' prices, and not via the common prices, because all shoppers' prices are below all common prices. Thus, when competition stiffens, firms start using the common prices even more for extracting profits from nonshoppers and, accordingly, raise the common prices. Higher common prices in turn slacken the competition for shoppers and firms raise their shoppers' prices. In combination, an increase in N can lead to an increase in all prices and an individual firm's profit. Figure 4 contrasts the results in my model for $\mu_s = \mu_n = \frac{1}{4}$ and the unit-demand version of Stahl (1989), i.e., $\mu_s = \mu_n = 0$. For these parameter values, in my model a firm's profit increases by about 2% if the market moves from four to five firms and by 15% if it moves to seven firms.

The proof derives strong sufficient conditions for a firm's profit to increase in the number of firms. One necessary condition is that the discrimination probability against shoppers, μ_s , is high enough.²¹ A high μ_s means that firms can use common prices mostly for extracting profits from nonshoppers. This ensures that the highest common price increases a lot if the number of firms increases, leading to a higher average price paid by a consumer. Firms extract this additional revenue and an individual firm's profit increases.

The crucial assumption for this result to hold is neither that the consumers

²⁰Neither Stahl (1989) nor I can prove his result analytically, but he verified it numerically.

²¹The inspection of equation (21) reveals that μ_s high enough is a necessary condition.

differ in their search costs nor that the firms can only discriminate between two groups of consumers. What matters is that, first, consumers differ in price-sensitivity and, second, that the prices set to at least some consumer groups are positively related in equilibrium. Then the mechanism that drives this paper’s results applies: the more price-sensitive consumers are on average offered lower prices than the less price-sensitive consumers. If the number of firms increases, firms have to lower the prices offered to price-sensitive consumers a lot to attract them. It can be profitable, instead, to focus on extracting profits from the price-insensitive consumers. This slackens competition, the price-sensitive consumers can also be offered higher prices and firms’ individual profits can increase.

This result implies that if firms can price discriminate based on search costs, an environment that looks more competitive, i.e., with many firms, may be much worse for consumers. A consumer protection policy that aims to increase competition by raising the number of sellers can, thus, instead backfire: hurt consumers and benefit firms. At the very least, many sellers is an unreliable measure of healthy competition online.

Empirical papers consistently find that price dispersion online increases in the number of firms (Nelson et al., 2007; Haynes and Thompson, 2008; Grewal et al., 2010; Böheim et al., 2021; Wang et al., 2021). However, the results on average prices are mixed: some papers find that the average price online increases in the number of firms (Haynes and Thompson, 2008; Grewal et al., 2010) while others find the opposite (Grewal et al., 2010; Wang et al., 2021). My model provides a joint explanation to the empirical observations that price dispersion and average price can increase in the number of firms online.

7 Extensions and discussion

I discuss various ways in which the model could be extended and argue that the model’s main results are unchanged.

7.1 Endogenous identification probabilities

The probabilities with which a shopper and a nonshopper are identified, μ_s and μ_n , can be endogenised in a variety of ways. Sections 7.1.1 and 7.1.2 discuss identification probabilities determined by consumers’ privacy choices and based on consumer behaviour. Effectively, each possibility (and other conceivable ones, e.g., based on consumer characteristics such as privacy-concern or computer-literacy types) simply pins down the identification probabilities at some concrete values.

7.1.1 Endogenous privacy choices

Suppose that a consumer can choose, at a privacy cost, to hide his type (search cost) from the firms rather than reveal his type. I argue that under reasonable assumptions on privacy costs, my main model's results hold.

In my model shoppers want to reveal and nonshoppers to hide their type because identified shoppers pay, on average, lower and nonshoppers higher prices than the others. How many consumers reveal their type in equilibrium depends on how privacy costs are modelled.

A natural assumption from the modelling viewpoint is that all consumers have an identical, potentially negative, privacy cost. Privacy cost is negative if verifiably revealing the type is costly and if privacy foregone in one market hurts the consumer in other markets. With identical privacy cost, the equilibrium depends on the size of the cost. If the privacy cost is positive, shoppers reveal their type and nonshoppers do not. Thus, a consumer's type can be inferred from his privacy choice. In equilibrium a nonshopper pays his cutoff price, whereas a shopper pays a price of zero. If the privacy cost is negative enough, no consumer reveals his type and in equilibrium firms offer only common prices.

Modelling consumers as having an identical privacy cost, however, ignores an important feature of online privacy: once individual data is revealed, a consumer cannot revoke the data easily, but a firm can easily share the data with agents on this or other markets. At least some consumers are aware of this so reasonable assumptions on privacy costs are that the costs differ across consumers and that some consumers have a negative privacy cost. In this case, in equilibrium a consumer's privacy choice depends on the comparison between his benefit and cost of retaining privacy. The benefits of retaining privacy are paying a lower price for a nonshopper and "paying" a negative privacy cost for both some shoppers and nonshoppers. The costs of retaining privacy are paying a higher price for a shopper and paying a positive privacy cost for both some shoppers and nonshoppers. As long as both some shoppers and some nonshoppers choose to retain and others to forego privacy, the same analysis applies as in my main model and its results continue to hold.

7.1.2 Behaviour-based identification probabilities

The main interpretation of the model that I suggest is that of web cookies. In line with this interpretation, suppose that the probability with which a consumer is identified increases in the number of firms that he visits.

In any equilibrium where shoppers' prices are on average lower than common

prices and common prices lower than nonshoppers' prices, a shopper has an incentive to be identified and a nonshopper does not. In other words, a shopper wants to visit as many firms as possible before buying (as a shopper is assumed to do if μ_s is exogenous). Conversely, a nonshopper wants to visit as few firms as possible before buying (as a nonshopper optimally does if μ_n is exogenous). Because nonshoppers are in general more profitable to serve than shoppers, firms collectively would want to offer with positive probability nonshoppers' or common prices that exceed ϕ_n in order to make nonshoppers visit more than a single firm and, thus, be able to identify them better. But any single firm would want to deviate from setting such prices because these prices are not accepted. Thus, the equilibrium characterisation would be the same as in the main model, except that the identification probabilities would be ordered $\mu_s > \mu_n$.

7.2 More consumer heterogeneity

In my main model, consumers differ only in their search cost, which can take two values. I discuss here how the model's results would be affected if consumers differed in computer literacy, if their search costs took on more than two values, or if their valuations were dispersed. The first extension is another way of pinning down the consumer identification probabilities. The two latter extensions lead to dispersed nonshoppers' prices and to a potential gap between the highest common and the highest nonshopper's price.

7.2.1 Computer literacy

Suppose that, in addition to search costs, consumers differ in their computer literacy. I show that my model can be extended to allow this second dimension of consumer heterogeneity without affecting its main results.

Assume that each consumer has a computer-literacy type: either he is computer-literate who can deal with the latest developments in computers and related technology (including smartphones) or computer-illiterate. The fraction of shoppers (nonshoppers) who are computer-literate is denoted by λ_{sl} (λ_{nl}) and who are computer-illiterate by λ_{si} (λ_{ni}) with $\lambda_{sl} + \lambda_{si} = 1$ and $\lambda_{nl} + \lambda_{ni} = 1$. I assume that the computer-literate consumers are better at searching online than computer-illiterate so $\lambda_{sl} > \lambda_{nl}$. Also, the computer-literate consumers understand the privacy risks associated with enabling cookies so they are more likely to switch off cookies: $\mu_l < \mu_i$, $\mu_l, \mu_i \in (0, 1)$. A firm identifies the search cost of a consumer that enables cookies.

Then the total probability that a firm identifies a visiting, say, shopper is

$\lambda_{sl}\mu_l + (1 - \lambda_{sl})\mu_i$, which is similar to the total probability of identifying a shopper in my main model, μ_s . The expression can easily be modified to allow for more computer-literacy types. The assumptions I made on the parameter values mean that the total probability of identifying a shopper is lower than that of identifying a nonshopper: $\mu_s < \mu_n$.

The analysis does not change much is that firms do not care about the consumers' computer literacy per se. Instead, firms care about how many other firms offer a shopper a discriminatory price and how many offer a common price. The set of price offers that a shopper receives is affected by his literacy type so the calculation of a firm's probability of selling to a shopper (see equation (5)) becomes more complex. In general, the selling probability has such a complex form that the equilibrium price distributions cannot be solved for. But the economic forces are the same in the extended and the main model. Unless consumers' computer-literacy and search types are perfectly correlated, all consumers are with positive probability offered both discriminatory and common prices. As a result, the incentives of firms to set the different prices are unchanged.

7.2.2 More dispersed search costs

Suppose that consumers' search costs are distributed according to $\alpha \sim G(\alpha)$ with support $[0, \bar{\alpha}]$. Let the probability with which type- α consumers are identified be $\mu_\alpha \in (0, 1)$ and assume that $G(0) > 0$.²²

Intuitively, symmetric equilibria in this version of the model, if they exist, must satisfy the following. All consumers with $\alpha > 0$ behave like nonshoppers in my main model: accept prices below their search-cost-dependent cutoff price ϕ_α and continue searching otherwise. Thus, an identified consumer with $\alpha > 0$ is offered a price p_α just equal to his cutoff price ϕ_α : the discriminatory price for every consumer with $\alpha > 0$ is a singleton.

Second, the discriminatory prices offered to consumers with $\alpha = 0$, the shoppers, are dispersed for the same reasons as in the main model. Any singleton price $p_{\alpha=0} > 0$ would be undercut by a competitor and $p_{\alpha=0} = 0$ is not optimal because it yields zero profits while a price $p_{\alpha=0} \in (0, \underline{p}_c)$ generates positive profits from consumers with $\alpha = 0$ that are offered a common price at all other firms.

Finally, common prices are dispersed in general, but it is difficult to say something more concrete about the distribution. The common prices are dispersed because a low common price is accepted by more consumers, but generates less

²²Stahl (1996) shows in a sequential search model without price discrimination that multiple pure-strategy symmetric equilibria exist at some price $p^* \leq v$ if $G(\alpha = 0) = 0$. This result carries over to my model: if $G(\alpha = 0)$, $p_\alpha = p_c = p^*$ in a symmetric equilibrium.

revenue per consumer than a high common price. However, the highest common price may now be lower than the highest discriminatory price $\phi_{\alpha_{max}}$. To see this, note that if $\mu_{\alpha_{max}}$ is almost one, a firm (almost) never competes for α_{max} -consumers with its common price. But as $\phi_{\alpha_{max}} > \phi_{\alpha}$ for all $\alpha \neq \alpha_{max}$, a common price equal to $\phi_{\alpha_{max}}$ would be accepted by (almost) no consumer that it is offered to so a firm could increase its profit by lowering its common price. In sum, the equilibrium characterisation would be similar to the main model's expect that nonshoppers' prices would be dispersed and the highest nonshoppers' price could exceed the highest common price.

7.2.3 Dispersed valuations

In the main model, I focus on identical consumer valuations, first, to distil the effects of price discrimination based on search cost and, second, to keep the model tractable. Suppose instead, that consumers differ in their valuations for the good. Different valuations generate more complicated search behaviour by nonshoppers and more dispersion in some prices, but should not overturn the model's results because the economic forces would remain the same.

In particular, in equilibrium shoppers as price-sensitive consumers are offered low discriminatory prices and a nonshoppers' discriminatory price still equals his cutoff price. But because nonshoppers have different valuations, their cutoff prices now differ so also the nonshoppers' prices are dispersed. Depending on the distribution of valuations, the highest nonshoppers' cutoff price can exceed the highest common price.

7.3 Correlated consumer identification

In the main model, the events that a consumer's type is identified by multiple firms are independent. Suppose, instead, that a consumer who is identified by one firm is more likely to be identified by another. An interpretation of positive correlation among the identification events is that consumers disable or enable third-party (rather than first-party) cookies. The analysis becomes much more complicated, but I argue that the underlying forces in the model do not change dramatically unless the correlation between the identification events is one.

Assume that the identification events of a single consumer are correlated: if a consumer is identified at the first firm he visits, he is more likely to be identified, rather than unidentified, at any subsequent firm he visits. A consumer knows whether he is identified at any firm. I discuss perfectly and imperfectly correlated identification of consumers in turn.

7.3.1 Perfectly correlated identification events

The simplest way of modelling positive correlation is that the identification events are perfectly correlated: with probability μ_s a shopper is identified by all firms and with probability $1 - \mu_s$ by no firm, and analogously for nonshoppers. An interpretation of this structure is that some consumers' data is in the database of an information intermediary and all firms have access to the data of the intermediary.

The head-on price competition for the identified shoppers drives their price to zero so $p_s = 0$. An identified nonshopper is offered p_n for sure at all firms he visits which drives their price to the valuation so $p_n = v$. The identified nonshoppers' option value of searching is effectively removed.

Firms use the common prices to compete over the unidentified consumers a la Stahl (1989) where the measure of shoppers who are offered a common price is $\lambda(1 - \mu_s)$ and the measure of nonshoppers is $(1 - \lambda)(1 - \mu_n)$. The highest common price \bar{p}_c^{corr} is the minimum of the consumers' valuation v and the unidentified nonshoppers' cutoff price: the price that leaves him just indifferent between stopping and continuing to search at other firms that cannot identify him. The highest common price becomes explicitly $\bar{p}_c^{corr} = \min \left\{ \alpha \left[1 - \int_0^1 (1 + \gamma^{corr} N y^{N-1})^{-1} dy \right]^{-1}, v \right\}$.

Since a firm must be indifferent between charging any common price with probability one, the expected profit that a firm earns in a symmetric equilibrium here is

$$\pi^{corr} = (1 - \lambda)\mu_n \frac{v}{N} + (1 - \lambda)(1 - \mu_n) \frac{\bar{p}_c^{corr}}{N},$$

because a firm that sets the highest common price sells to no shoppers. If \bar{p}_c^{corr} is interior, the expected profit π^{corr} increases in v . If \bar{p}_c^{corr} has a boundary solution (i.e., if v is low), the profit becomes $\pi^{corr} = (1 - \lambda) \frac{v}{N}$.

In the main model a firm's expected equilibrium profit is

$$\pi = (1 - \lambda) \frac{\bar{p}_c}{N} + \lambda \mu_s (1 - \mu_s)^{N-1} \frac{\bar{p}_c}{1 + \gamma N}.$$

For low enough consumers' valuation v , both the highest common prices \bar{p}_c^{corr} and \bar{p}_c have a boundary solution: $\bar{p}_c^{corr} = \bar{p}_c = v$. In this case, firms earn higher profits when the identification events are independent rather than perfectly correlated. In other words, firms prefer to track consumers independently rather than buy data from a data intermediary.

If the consumer's valuation v is high enough, the highest common prices \bar{p}_c and \bar{p}_c^{corr} are interior, independent of v , and satisfy $v > \bar{p}_c > \bar{p}_c^{corr}$. For high enough v , thus, firms earn lower profits when the identification events are independent rather than perfectly correlated.

In sum, firms prefer independent identification events if the consumers' valuation is low and perfectly correlated identification events otherwise. In other words, firms profit from having access to the data of a single information intermediary only if consumers value the product highly. Firms benefit from removing the option value of searching from identified nonshoppers only if the consumers' valuation of the good is high.

7.3.2 Imperfectly correlated identification events

If the positive correlation is imperfect, the analysis becomes more cumbersome, but the economic forces remain the same as in the main model. First, consider the nonshoppers. Since the identification events of a nonshopper are correlated, both his continuation value and cutoff price depend on whether he is or is not identified by the firm he visits now.²³ His cutoff price after being identified is (weakly) higher than after being unidentified because the firms' optimal nonshoppers' price is always the highest acceptable price. For a firm, this is the only consideration that matters about the correlation between nonshoppers' identification events because a firm still has a temporary monopoly power over a visiting nonshopper. The wedge between the cutoff prices after being identified versus unidentified creates a wedge between the nonshoppers' price and the highest common price: in equilibrium the nonshoppers' price is above the highest common price.

Now consider the shoppers and the expression for the probability that an identified shopper buys at a firm, equation (5) for $p^i = p_s^i$. As compared to independence, a positive correlation among firms' identification events means that it becomes more likely that very many or very few firms identify a particular shopper (and less likely that an intermediate number of firms identify the shopper). In the first line of equation (5) for $p^i = p_s^i$, the probability with which exactly k firms identify a shopper would depend on k in a more complex manner than k appearing just in the exponent. In general, the probabilities associated with large and small k would increase, and associated with intermediate k would decrease. For even the simplest way of modelling imperfect correlation, the expression can no longer be solved for the equilibrium price distributions.

But the same forces are at play in this extended model as in my main model. In particular, a firm that can identify a shopper is interested in how many other firms do so, too: this determines how fiercely the firm competes for shoppers with its shoppers' price. With positively correlated identification events of shoppers, firms

²³Whether only the last or the entire sequence of past identification events matters depends on the correlation structure. The argument here holds if only the identification event at the current firm matters for the likelihood of being identified at the next firm.

compete more fiercely for shoppers using their shoppers' prices so I expect that the shoppers' prices decrease. The effect on common prices depends on whether only shoppers' or both shoppers' and nonshoppers' identification events are correlated and on features of the correlation structure. As long as common prices serve both shoppers and nonshoppers, the common prices should on average exceed shoppers' prices. In sum, I expect the equilibrium to still feature low and dispersed shoppers' prices, higher dispersed common prices, and a singleton nonshoppers' price.

8 Conclusion

Advances in ICT have made tracking people's behaviour online, and both storing and analysing the resulting data, considerably cheaper. As a result, price discrimination based on people's online behaviour has become feasible. This paper analyses a model where firms can infer people's search cost (from their online behaviour) and probabilistically price discriminate based on this information. Price discrimination can be seen as probabilistic if some consumers disable cookies or if firms cannot track consumers perfectly. Regulations lower the probability of price discrimination online if they facilitate the disabling of cookies or restrict tracking.

According to my model, in general consumers lose and firms benefit from more likely price discrimination. In the presence of price discrimination, not only does the industry profit increase in the number of firms, but so does an individual firm's profit for certain parameter values. In other words, a firm can have strict incentives to attract competitors to the market and also benefits from splitting and selling as multiple independent entities. This type of spurious competition hurts consumers and the potential for this harm is especially large in online markets where inferring consumers' search costs is feasible. The crucial assumption for this result to hold is neither that the consumers differ in their search costs nor that the firms can only discriminate between two groups of consumers.

The potential to personalise prices seems to be ever increasing, while detecting personalised prices remains hard. On the one hand, developments in AI and in firms' capacity to analyse big data facilitate price discrimination.²⁴ AI techniques can be used to solve complex optimisation problems and adjust prices automatically and dynamically. Big data can be used to target prices very precisely. On the other, detecting personalised prices online remains hard. Empirical studies, such as Ipsos et al. (2018) and Ennis and Lam (2021), that document behaviour-based price discrimination online have overcome the difficulty of detecting personalised

²⁴Bourreau and de Streel (2018), Ipsos et al. (2018), Ennis and Lam (2021)

prices in clever ways, but each involves problems. Simulated consumer profiles generally cannot possess long and varied histories of online behaviour. Using the profiles of real people may cause ethical problems because the studies may affect the prices the people face later. Detecting personalised prices is harder now than even a decade ago because more people use mobile phones, where the browsers' default is to allow tracking, to browse the internet.²⁵ As a result, price discrimination online promises to be an important research topic for the foreseeable future.

A Appendix

Here are the proofs omitted from the paper.

A.1 Equilibrium

Proof of Proposition 1. I derive the equilibrium in steps. Steps 1-4 derive some necessary conditions on the equilibrium price distributions. For Steps 1-4, I assume that $\underline{p}_s \leq \underline{p}_c$ and $\bar{p}_s \leq \bar{p}_c$ (which I verify in Step 6). Step 5 argues that the Proposition describes an equilibrium. Step 6 shows that the equilibrium is the unique symmetric equilibrium. The derivation is for all $\mu_s, \mu_n \in (0, 1)$

Throughout, consumers' passive beliefs imply that a nonshopper accepts any deviating price offer that is below ϕ_n and continues to search after any deviating price offer that exceeds ϕ_n . A shopper accepts any deviating price offer if it is the lowest among the price offers he receives.

Recall first that I showed in Section 4 that each firm i optimally sets its prices p_n^i , p_s^i and p_c^i independently of each other. Thus, I can separately derive equilibrium price distributions F_n , F_s and F_c .

Step 1: The equilibrium distribution of nonshoppers' prices $F_n(p_n)$ and the highest common price \bar{p}_c . If $\bar{p}_c \geq \bar{p}_s$ (verified in Step 6), the optimal nonshoppers' price in all symmetric equilibria is a singleton $p_n = \phi_n$ and the highest common price satisfies $\bar{p}_c \leq \phi_n$.

Suppose, instead, that the highest equilibrium nonshoppers' or common price exceeds ϕ_n : $\max\{\bar{p}_n, \bar{p}_c\} > \phi_n$. Assume first also that $\bar{p}_n \geq \bar{p}_c$. Then firm i would never want to set $p_n^i = \bar{p}_n$: all nonshoppers who visit firm i and i identifies would also visit another firm, get a price offer below \bar{p}_n there, and thus never return to i . Firm i would be better off deviating and setting $p_n^i = \phi_n$. Assume now that

²⁵In EU27, 16% of people used a mobile phone to access the internet in 2011 and 71% in 2019 Eurostat (2021). The default browser neither on Android phones nor iPhones has private browsing as the default choice in 2022.

$\bar{p}_c \geq \bar{p}_n$ and that the highest shoppers' price is weakly below the highest common price ($\bar{p}_c \geq \bar{p}_s$). Then firm i would never want to set $p_c^i = \bar{p}_c$: all nonshoppers who visit firm i would also visit another firm, none would return to i , and all shoppers would buy at another firm. Firm i would be better off deviating and setting a lower p_c^i . Thus, $\bar{p}_n \leq \phi_n$ and $\bar{p}_c \leq \phi_n$ in all symmetric equilibria.

So for firm i 's equilibrium nonshoppers' price, in the interim profit from an identified nonshopper, equation (2), only the first line is relevant. Since this increases in p_n^i , in all symmetric equilibria the optimal nonshoppers' price is $p_n = \phi_n$ and $F_n(p_n)$ is degenerate. The firms simply exercise their temporary monopoly power over the visiting identified nonshoppers.

Step 2: Explicit forms for the probability that a shopper buys at i at price p^i and the ex ante expected profits. Since firms $j \neq i$ offer only prices weakly below the nonshoppers' cutoff price in symmetric equilibria, no nonshopper that visited firm $j \neq i$ first visits also firm i : no nonshopper returns to firm i so $P(n \text{ returns to } i | p^i) = 0$ for $p^i = p_n^i, p_c^i$ and the total probability that a nonshopper visits firm i is $P(\text{visits } i | n) = \frac{1}{N}$. In symmetric equilibria, thus, firm i 's posterior belief about the type of a visiting unidentified consumer is

$$\frac{\nu(s|\text{unid})}{1 - \nu(s|\text{unid})} = \frac{\lambda}{1 - \lambda} \frac{1 - \mu_s}{\frac{1}{N} - \mu_n},$$

and the interim profit from an unidentified consumer, equation (4), simplifies considerably.

To complete the derivation of interim profits, let us turn to the probability that a shopper who is offered price p^i at firm i buys at i , $D_s^i(p^i)$. Because all firms $j \neq i$ use the same independent price distributions F_s and F_c for their shoppers' and common prices respectively in a symmetric equilibrium, $D_s^i(p^i)$ can be written explicitly as

$$\begin{aligned} D_s^i(p^i) &= \sum_{k=0}^{N-1} \binom{N-1}{k} \mu_s^k (1 - F_s(p^i))^k (1 - \mu_s)^{N-1-k} (1 - F_c(p^i))^{N-1-k} \\ &= [\mu_s(1 - F_s(p^i)) + (1 - \mu_s)(1 - F_c(p^i))]^{N-1}, \end{aligned} \quad (5)$$

if p^i is in the support of both F_s and F_c . Firm i behaves as if it competes for shoppers against a mixture distribution where a competing price comes from F_s with probability μ_s and from F_c with probability $1 - \mu_s$.

The expected demand changes a bit if firm i offers a shopper a price p^i that is in the support of, say, F_s but not in that of F_c . Intuitively, the equilibrium

shoppers' prices should be weakly lower than the common prices, i.e., $\underline{p}_s \leq \underline{p}_c$ and $\bar{p}_s \leq \bar{p}_c$. This is because a firm's common price competes for both nonshoppers and shoppers, but its shoppers' price competes only for shoppers and shoppers are more price-sensitive than nonshoppers. If $\underline{p}_s \leq \underline{p}_c$ and $\bar{p}_s \leq \bar{p}_c$, then the probability that a shopper who gets price offer p^i at firm i buys at i , $D_s^i(p^i)$, has at most three parts in equilibrium:

$$D_s^i(p^i) = \begin{cases} [\mu_s(1 - F_s(p^i)) + 1 - \mu_s]^{N-1} & \text{for } p^i \in [\underline{p}_s, \underline{p}_c), \\ [\mu_s(1 - F_s(p^i)) + (1 - \mu_s)(1 - F_c(p^i))]^{N-1} & \text{for } p^i \in [\underline{p}_c, \bar{p}_s), \\ [(1 - \mu_s)(1 - F_c(p^i))]^{N-1} & \text{for } p^i \in [\bar{p}_s, \bar{p}_c]. \end{cases} \quad (6)$$

In total, firm i 's ex ante profits when setting prices (p_c^i, p_s^i, p_n^i) are

$$\pi^i = \frac{1 - \lambda}{N} \mu_n \pi_n^i(p_n^i) + \lambda \mu_s \pi_s^i(p_s^i) + \left[\frac{1 - \lambda}{N} (1 - \mu_n) + \lambda (1 - \mu_s) \right] \pi_c^i(p_c^i). \quad (7)$$

The amount of nonshoppers that visit firm i is $\frac{1-\lambda}{N}$. Each is offered the nonshoppers' price with probability μ_n and the common price with probability $1 - \mu_n$, yielding interim profits $\pi_n^i(p_n^i)$ and $\pi_c^i(p_c^i)$. The amount of shoppers that visit firm i is λ . Each is offered the shoppers' price with probability μ_s and the common price with probability $1 - \mu_s$, yielding interim profits $\pi_s^i(p_s^i)$ and $\pi_c^i(p_c^i)$.

Step 3: Necessary conditions on an equilibrium distribution of shoppers' prices $F_s(p_s)$. For this step, I also assume that $\underline{p}_c > 0$ (verified in Step 4). Lemmas 1-3, that are quite standard, derive some necessary conditions on F_s .

Lemma 1. *The shoppers' discriminatory prices p_s are drawn from a distribution $F_s(p_s)$ that has no mass points on any $p_s \geq 0$ if $\underline{p}_c > 0$.*

Proof of Lemma 1. Suppose first that $\underline{p}_s > 0$ and that all firms use F_s that has a mass point of size \hat{f} at some $p_s = p$. Then firm i can profitably deviate by moving the mass point in its distribution F_s^i to $p_s^i = p - \varepsilon$. The deviation is profitable because firm i 's probability of serving a shopper jumps up by a discrete amount, whereas its revenue drops only by a bit: now firm i gets a shopper also if all other firms offer the shopper the shoppers' price p , with probability $\mu_s^{N-1} \hat{f}^{N-1}$.

Now suppose that firms use F_s that has a mass point on $p_s = 0$ so that in equilibrium $\pi_s^i(p_s^i) = 0$. But then firm i can profitably deviate by setting the shoppers' price $p_s^i \in (0, \underline{p}_c)$. Firm i serves all shoppers who only it identifies and earns positive expected profit of $\lambda p_s^i (1 - \mu_s)^{N-1}$ from these shoppers. \square

Lemma 2. *The support of $F_s(p_s)$ is an interval $[\underline{p}_s, \bar{p}_s]$.*

Proof of Lemma 2. Suppose instead that all firms use F_s that puts no weight on some interval $[p_1, p_2]$. Then firm i can profitably move mass from the interval $(p_1 - \varepsilon, p_1)$ to p_2 : i does not lose demand from shoppers, but increases the expected price from them by a discrete amount. \square

If a firm mixes over shoppers' prices in equilibrium, it must be indifferent between setting each shoppers' price $p_s \in [\underline{p}_s, \bar{p}_s]$ with probability one. If firm i sets $p_s^i = \bar{p}_s$, its expected equilibrium profit from shoppers who i identifies is

$$k := \lambda \pi_s^i(\bar{p}_s) = \lambda(1 - \mu_s)^{N-1}(1 - F_c(\bar{p}_s))^{N-1} \bar{p}_s. \quad (8)$$

In equilibrium k is a positive constant if $\underline{p}_c > 0$ and $\lambda \pi_s^i(p_s^i) = k$ must hold for all shoppers' prices in the support of F_s . Using the equal-profit condition and equation (6), we get that an equilibrium F_s must satisfy

$$\mu_s(1 - F_s(p_s)) = \begin{cases} \left(\frac{k}{\lambda p_s}\right)^{\frac{1}{N-1}} - (1 - \mu_s) & \text{for } p_s \in [\underline{p}_s, \underline{p}_c], \\ \left(\frac{k}{\lambda p_s}\right)^{\frac{1}{N-1}} - (1 - \mu_s)(1 - F_c(p_s)) & \text{for } p_s \in [\underline{p}_c, \bar{p}_s]. \end{cases} \quad (9)$$

Lemma 3. $\bar{p}_s < \bar{p}_c$ and $\underline{p}_s > 0$ if $\underline{p}_c > 0$.

Proof of Lemma 3. First, I show that $\bar{p}_s < \bar{p}_c$ if $\underline{p}_s > 0$. At $p_s = \bar{p}_s$, a firm's expected profit from identified shoppers is given by equation (8). If $\bar{p}_s = \bar{p}_c$, then $F_c(\bar{p}_s) = 1$ and the equilibrium expected profit from shoppers is zero. But then firm i has a profitable deviation to $p_s = \underline{p}_s - \varepsilon > 0$ (see the proof of Lemma 1).

Second, I show that $\underline{p}_s > 0$. Since I have not shown that there is no mass point on \underline{p}_s , let $F_s(\underline{p}_s) = \hat{f} \in [0, 1)$ and use equation (9) to rewrite \underline{p}_s as $\underline{p}_s = k\lambda^{-1}[\mu_s(1 - \hat{f}) + (1 - \mu_s)]^{-(N-1)}$. This is positive because $k > 0$ if $\underline{p}_c > 0$. \square

So far we know that in equilibrium, shoppers' prices are drawn from the mixed-strategy distribution $F_s(p_s)$ with support $[\underline{p}_s, \bar{p}_s]$ where $0 < \underline{p}_s < \bar{p}_s < \bar{p}_c$ if $\underline{p}_c > 0$.

Step 4: Necessary conditions on an equilibrium distribution of common prices $F_c(p_c)$. I argued in Step 1 that in symmetric equilibria the highest common price is equal to the nonshoppers' cutoff price ϕ_n . Using equations (6), (7), and (9), I can rewrite firm i 's expected profit when setting a common price p_c^i as

$$\pi^i(p_c^i) = \begin{cases} \frac{1-\lambda}{N}[(1 - \mu_n)p_c^i + \mu_n\phi_n] + k & \text{for } p_c^i \in [\underline{p}_c, \bar{p}_s], \\ \frac{1-\lambda}{N}[(1 - \mu_n)p_c^i + \mu_n\phi_n] + (1 - \mu_s)\lambda[(1 - \mu_s)(1 - F_c(p_c^i))]^{N-1}p_c^i + \mu_s k & \text{for } p_c^i \in [\bar{p}_s, \bar{p}_c], \end{cases} \quad (10)$$

because, when considering which common price to set, both ϕ_n and k are constants from firm i 's viewpoint. But then the first part of equation (10) increases in p_c^i , which cannot describe the expected profit from mixing over $p_c^i \in [\underline{p}_c, \bar{p}_s]$: that needs to be constant in p_c^i . As a result, a nondegenerate common price distribution F_c has support $[\underline{p}_c, \bar{p}_c]$, with $\underline{p}_c \geq \bar{p}_s$ and a possible mass point at \underline{p}_c if $\underline{p}_c = \bar{p}_s$.

Lemma 4. $\underline{p}_c > 0$.

Proof of Lemma 4. Suppose the contrary: $\underline{p}_c = 0$. Then the equilibrium profits of a firm that sets $p_c = \underline{p}_c$ are zero from unidentified (and identified) shoppers and from unidentified nonshoppers. But then firm i can profitably deviate to $p_c^i \in (0, \alpha)$: it does not gain any custom from shoppers, but gets a positive profit of $\frac{1-\lambda}{N}(1-\mu_n)p_c^i$ from unidentified nonshoppers that visit i first. \square

Lemma 5. *The common prices p_c are drawn from a distribution $F_c(p_c)$ that has no mass points on any p_c .*

Proof of Lemma 5. The proof is analogous to that of Lemma 1, except that a deviating firm increases its profit if all other firms offer the shopper the common price p , which happens with probability $(1-\mu_s)^{N-1}\hat{f}^{N-1}$. \square

Lemma 6. $\bar{p}_s = \underline{p}_c$.

Proof of Lemma 6. A firm's interim profit per identified shopper when setting $p_s = \bar{p}_s$ is $\pi_s(\bar{p}_s) = \bar{p}_s(1-\mu_s)^{N-1}$ as $\bar{p}_s \leq \underline{p}_c$. Suppose that $\bar{p}_s < \underline{p}_c$. Then the firm can increase its profit by setting $p_s = \bar{p}_s + \varepsilon < \underline{p}_c$ instead of \bar{p}_s : its probability of serving a shopper is the same as when setting \bar{p}_s , but its revenue increases. \square

This completes the derivation of a potential equilibrium shoppers' price distribution F_s that I summarise in

Corollary 1. *An equilibrium distribution of shoppers' prices $F_s(p_s)$ must satisfy*

$$F_s(p_s) = \mu_s^{-1} \left[1 - (1 - \mu_s) \left(\frac{\bar{p}_s}{p_s} \right)^{\frac{1}{N-1}} \right],$$

for $p_s \in [\underline{p}_s, \bar{p}_s]$ where $\underline{p}_s = (1 - \mu_s)^{N-1}\bar{p}_s$ and $\bar{p}_s = \underline{p}_c$.

Proof. The distribution is derived from equation (9), using the fact that $\bar{p}_s = \underline{p}_c$ and the definition of k : $k = \lambda(1 - \mu_s)^{N-1}\bar{p}_s$. Using $F_s(\underline{p}_s) = 0$, we get \underline{p}_s . \square

Let us return to the distribution of common prices. Gaps in the support of F_c can be ruled out in an analogous manner as for F_s (see Lemma 2). Since firm i must be indifferent across setting any common price in the support of F_c

with probability one, I get a necessary condition on an equilibrium distribution F_c from solving $\pi^i(\bar{p}_c) = \pi^i(p_c)$ (see the relevant, second, part of equation (10)). A potential equilibrium distribution of common prices F_c is summarised in

Corollary 2. *An equilibrium distribution of common prices $F_c(p_c)$ must satisfy*

$$F_c(p_c) = 1 - \left(\frac{\bar{p}_c - p_c}{\gamma N p_c} \right)^{\frac{1}{N-1}}, \quad (11)$$

for $p_c \in [\underline{p}_c, \bar{p}_c]$ where $\bar{p}_c = \min \left\{ \frac{\alpha}{1-\mu_n} \left[1 - \int_0^1 (1 + \gamma N y^{N-1})^{-1} dy \right]^{-1}, v \right\}$, $\underline{p}_c = \frac{\bar{p}_c}{1+\gamma N}$, and $\gamma := \frac{\lambda(1-\mu_s)^N}{(1-\lambda)(1-\mu_n)}$.

Proof. The lowest common price \underline{p}_c can be solved from $F_c(\underline{p}_c) = 0$ which gives

$$\underline{p}_c = \bar{p}_c(1 + \gamma N)^{-1}. \quad (12)$$

I showed in Step 1 that $\bar{p}_c \leq \phi_n$. Suppose that $\bar{p}_c < \phi_n$. Then a firm that sets $p_c = \bar{p}_c$ could increase its revenue from unidentified nonshoppers (and would not lose demand from unidentified shoppers, which is zero at $p_c = \bar{p}_c$) by setting $p_c \in (\bar{p}_c, \phi_n)$. Thus, in equilibrium $\bar{p}_c = \phi_n$.

From Section 3 we know that ϕ_n satisfies equation (2) if the solution to (2) is below v . If the nonshopper's best offer so far is $p = \phi_n$, the benefit of searching another firm is

$$B(\phi_n) = (1 - \mu_n) \int_{\underline{p}_c}^{\phi_n} (p - p_c) dF_c(p_c).$$

because in equilibrium $\phi_n = p_n$. Since $\bar{p}_c = \phi_n$, integration gives $B(\phi_n) = (1 - \mu_n)[\phi_n - \mathbb{E}_c(p_c)]$. If the solution to $B(\phi_n) = \alpha$ is below the valuation v , then ϕ_n solves

$$\phi_n = \mathbb{E}_c[p_c] + \frac{\alpha}{1 - \mu_n}.$$

I use the methods in Janssen et al. (2005) to solve for interior ϕ_n explicitly:

$$\phi_n = \frac{\alpha}{1 - \mu_n} \left[1 - \int_0^1 (1 + \gamma N y^{N-1})^{-1} dy \right]^{-1}. \quad (13)$$

If the solution to the RHS of (13) is below v , it gives both the cutoff price ϕ_n and the highest common price \bar{p}_c . If the solution exceeds v , then $\phi_n = \bar{p}_c = v$. \square

Step 5. Proposition 1 describes an equilibrium. By construction, I have shown that if other firms mix over p_c and p_s using F_c and F_s as described in Proposition 1, then a single firm i is indifferent between setting with probability

one all $p_c \in [\underline{p}_c, \bar{p}_c]$ and all $p_s \in [\underline{p}_s, \bar{p}_s]$. I show that setting a common price outside $[\underline{p}_c, \bar{p}_c]$ or a shoppers' price outside $[\underline{p}_s, \bar{p}_s]$ yields a lower profit to firm i .

Firm i 's optimally sets its shoppers' and common prices independently. This also holds when firm i considers deviating so we can consider separately deviations from the proposed equilibrium common and shoppers' prices. Suppose that all other firms play according to the Proposition and consider the following deviations by firm i . A deviation profit is marked with a tilde.

First, deviating to a common price $p_c^i \in [\underline{p}_s, \underline{p}_c)$ is dominated by setting the lowest equilibrium common price. The deviation price reduces revenue and generates no larger demand. In particular, the deviation yields firm i expected profit

$$\begin{aligned}\tilde{\pi}^i(p_c^i) &= \frac{1-\lambda}{N}(1-\mu_n)p_c^i + (1-\mu_s)\lambda p_c^i [1-\mu_s + \mu_s(1-F_s(p_c^i))]^{N-1} + \frac{1-\lambda}{N}\mu_n\pi_n^i + \lambda\mu_s\pi_s^i \\ &= \frac{1-\lambda}{N}(1-\mu_n)p_c^i + (1-\mu_s)\lambda\bar{p}_s(1-\mu_s)^{N-1} + \frac{1-\lambda}{N}\mu_n\pi_n^i + \lambda\mu_s\pi_s^i,\end{aligned}$$

where the second line follows from plugging in F_s from the Proposition. Since

$$\tilde{\pi}^i(p_c^i) < \frac{1-\lambda}{N}(1-\mu_n)\underline{p}_c + (1-\mu_s)\lambda\underline{p}_c(1-\mu_s)^{N-1} + \frac{1-\lambda}{N}\mu_n\pi_n^i + \lambda\mu_s\pi_s^i = \pi^i(\underline{p}_c),$$

this deviation is unprofitable. By an analogous argument, a deviation to $p_c^i < \underline{p}_s$ is unprofitable.

Second, deviating to a common price $p_c^i > \bar{p}_c$ is dominated by setting the highest equilibrium common price because the deviation only reduces demand from unidentified nonshoppers. Altogether, firm i finds it unprofitable to set a common price outside $[\underline{p}_c, \bar{p}_c]$ if other firms mix over p_c and p_s using respectively F_c and F_s as described in Proposition 1.

Third, deviating to a shoppers' price $p_s^i < \underline{p}_s$ is dominated by setting the shoppers' price \underline{p}_s because the deviation only reduces revenue from shoppers.

Finally, deviating to a shoppers' price $p_s^i \in (\bar{p}_s, \bar{p}_c]$ is dominated by setting the shoppers' price \bar{p}_s because the deviation only reduces expected demand from shoppers. Thus, this deviation is also unprofitable. Altogether, firm i finds it unprofitable to set a shoppers' price outside $[\underline{p}_s, \bar{p}_s]$ if other firms mix over p_c and p_s using respectively F_c and F_s as described in Proposition 1. Thus, the strategies in the Proposition constitute an equilibrium.

Step 6. The equilibrium is the unique symmetric equilibrium. In Steps 1-4, I assumed that $\underline{p}_s \leq \underline{p}_c$ and $\bar{p}_s \leq \bar{p}_c$. To prove uniqueness, I rule out other configurations.

First suppose that $\underline{p}_s > \underline{p}_c$ in equilibrium. I show that then a firm has an incentive to deviate. First note that if $\underline{p}_s > \underline{p}_c$, then $\underline{p}_s > 0$ as $\underline{p}_c \geq 0$ must hold in equilibrium (the argument in Lemma 4 applies). As a result, Lemma 3 applies and $\bar{p}_s < \bar{p}_c$. Then the probability that an unidentified shopper buys from firm i , equation (6), has an additional part if firm i sets a common price $p_c^i \in [\underline{p}_c, \underline{p}_s]$:

$$D_s^i(p_c^i) = [\mu_s + (1 - \mu_s)(1 - F_c(p_c^i))]^{N-1} \quad \text{for } p_c^i \in [\underline{p}_c, \underline{p}_s].$$

As a result, firm i 's expected profit from setting $p_c^i = p$ also has an additional part: instead of equation (10), it reads

$$\pi^i(p) = \begin{cases} \Pi_n(p) + \mu_s k + (1 - \mu_s)\lambda[\mu_s + (1 - \mu_s)(1 - F_c(p))]^{N-1}p & \text{for } p \in [\underline{p}_c, \underline{p}_s), \\ \Pi_n(p) + k & \text{for } p \in [\underline{p}_s, \bar{p}_s), \\ \Pi_n(p) + \mu_s k + (1 - \mu_s)\lambda[(1 - \mu_s)(1 - F_c(p))]^{N-1}p & \text{for } p \in [\bar{p}_s, \bar{p}_c], \end{cases} \quad (14)$$

where I have denoted $\Pi_n(p) := \frac{1-\lambda}{N}[(1 - \mu_n)p + \mu_n\phi_n]$ for brevity. Equation (14) is increasing in p in the interval $p \in [\underline{p}_s, \bar{p}_s)$ so the only possibility for an equilibrium F_c is that it puts no mass on $p_c \in [\underline{p}_s, \bar{p}_s)$, and possibly a mass point on $p_c = \bar{p}_s$.

If F_c assigns no mass on $p_c \in [\underline{p}_s, \bar{p}_s)$, then the probability that an identified shopper buys from firm i 's if i sets $p_s^i \in [\underline{p}_s, \bar{p}_s]$ is

$$D_s^i(p_s^i) = [\mu_s(1 - F_s(p_s^i)) + (1 - \mu_s)(1 - F_c(\underline{p}_s))]^{N-1} \quad \text{for } p_s^i \in [\underline{p}_s, \bar{p}_s].$$

But then F_c cannot put a mass point on $p_c = \bar{p}_s$: otherwise, the probability that an identified shopper buys from firm i would have a jump when setting $p_s^i = \bar{p}_s$ instead of $p_s^i = \bar{p}_s - \varepsilon$, which cannot be in equilibrium (by a similar argument as in Lemma 1).

I argue that F_c assigns no mass on $p_c \in [\underline{p}_s, \bar{p}_s)$, then firm i has an incentive to deviate even without a mass point on $p_c = \bar{p}_s$. In particular, instead of setting $p_c^i \in [\underline{p}_s - \varepsilon, \underline{p}_s)$ with positive probability, firm i can increase its profit by assigning that probability to $p_c^i = \bar{p}_s$. The firm's demand would not drop, but its expected profit would increase by a discrete amount. Altogether, in equilibrium $\underline{p}_s \leq \underline{p}_c$.

Now suppose that $\bar{p}_s > \bar{p}_c$. I show that then a firm has an incentive to deviate. If $\bar{p}_s > \bar{p}_c$ and firm i sets $p_s^i > \bar{p}_c$, an identified shopper buys from i only if he also draws a shoppers' price from all the other firms:

$$D_s^i(p_s^i) = [\mu_s(1 - F_s(p_s^i))]^{N-1} \quad \text{for } p_s^i \in [\bar{p}_c, \bar{p}_s].$$

Firm i 's interim profit from the identified shopper is $\pi_s(p_s^i) = p_s[\mu_s(1 - F_s(p_s^i))]^{N-1}$.

By an analogous argument as in Lemma 1, p_s are dispersed in equilibrium so $\lambda\pi_s(p_s^i) = k$ must hold for all p_s^i in the support of F_s and some constant k . Thus, an equilibrium F_s must satisfy $\mu_s(1 - F_s(p_s)) = \left(\frac{k}{\lambda p_s}\right)^{\frac{1}{N-1}}$ for $p_s \geq \bar{p}_c$. Then setting a common price $p_c \in [\underline{p}_c, \bar{p}_c]$ would yield a profit $\pi(p_c)$ as described in the first line of equation (10). Since that increases in p_c , the only possible equilibrium F_c is degenerate at $p_c = \bar{p}_c$. But mass points on any $p_c > 0$ are ruled out by a similar argument as in Lemma 5 and $\bar{p}_c > 0$ must hold in equilibrium (by an analogous argument as in Lemma 4). Thus, in equilibrium $\bar{p}_s \leq \bar{p}_c$.

In sum, Proposition 1 describes the unique symmetric equilibrium. \square

A.2 Comparative statics

Proof of Proposition 2. Let us set the probabilities of price discrimination to be the same: $\mu_n = \mu_s =: \mu$. Then the boundary prices are $\underline{p}_c = \frac{\bar{p}_c}{1+\gamma N}$, $\bar{p}_c = \phi_n = \min\left\{\frac{\alpha}{1-\mu}(1-K)^{-1}, v\right\}$, and $\underline{p}_s = (1-\mu)^{N-1}\underline{p}_c$ where $K := \int_0^1 (1+\gamma N y^{N-1})^{-1} dy$ and $\gamma := \frac{\lambda(1-\mu)^{N-1}}{1-\lambda}$. The expected equilibrium profit is $\pi = \frac{1-\lambda}{N}\bar{p}_c + \mu\lambda\underline{p}_s$. The comparative statics with respect to μ are the following.

(i) I first cover the case when ϕ_n has an interior solution. I show in (a) that this holds for small μ : iff $\mu < \bar{\mu}$.

(a) \underline{p}_c and \bar{p}_c increase in μ . If \bar{p}_c is interior, by directly taking derivatives we get

$$\frac{\partial \bar{p}_c}{\partial \mu} = \bar{p}_c(1-\mu)^{-1} + \bar{p}_c(1-K)^{-1} \frac{\partial K}{\partial \gamma} \frac{\partial \gamma}{\partial \mu},$$

where $\frac{\partial K}{\partial \gamma} = -\int_0^1 (1+\gamma N y^{N-1})^{-2} N y^{N-1} dy < 0$ and $\frac{\partial \gamma}{\partial \mu} = -\frac{(N-1)\gamma}{1-\mu} < 0$. Thus, if \bar{p}_c is interior, then $\frac{\partial \bar{p}_c}{\partial \mu} > 0$. The interior solution to \bar{p}_c tends to $+\infty$ as $\mu \rightarrow 1$ so we know that ϕ_n is interior iff $\mu < \bar{\mu}$ for some $\bar{\mu} < 1$.

Also \underline{p}_c increases in μ because \underline{p}_c increases in \bar{p}_c and decreases in γ .

(b) \underline{p}_s increases in μ . The derivative with respect to \underline{p}_s is

$$\frac{\partial \underline{p}_s}{\partial \mu} = -(N-1) \frac{(1-\mu)^{N-2}}{1+\gamma N} \bar{p}_c + \frac{(1-\mu)^{N-1}}{1+\gamma N} \frac{\partial \bar{p}_c}{\partial \mu} - \frac{(1-\mu)^{N-1}}{(1+\gamma N)^2} \bar{p}_c N \frac{\partial \gamma}{\partial \mu}.$$

Plugging in the derivatives $\frac{\partial \gamma}{\partial \mu}$ and $\frac{\partial \bar{p}_c}{\partial \mu}$ from above, dividing the resulting equation with $\frac{(1-\mu)^{N-2}\gamma N}{\bar{p}_c(1+\gamma N)^2(1-K)}$ and collecting terms gives

$$\frac{\partial \underline{p}_s}{\partial \mu} \propto \int_0^1 \frac{y^{N-1}[\gamma N^2(1-y^{N-1}) + \gamma N y^{N-1}(2+\gamma N) + 1]}{(1+\gamma N y^{N-1})^2} dy > 0.$$

(c) π increases in μ . We have $\frac{\partial \pi}{\partial \mu} >$ because $\frac{\partial \bar{p}_c}{\partial \mu} > 0$ and $\frac{\partial \underline{p}_s}{\partial \mu} > 0$.

(ii) I now cover the case when ϕ_n has a boundary solution, i.e., $\mu \geq \bar{\mu}$. If $\phi_n = v$, then $\bar{p}_c = v$, $\underline{p}_c = v(1 + \gamma N)^{-1}$, $\underline{p}_s = v(1 - \mu)^{N-1}(1 + \gamma N)^{-1}$, and $\pi = \frac{1-\lambda}{N}v + \mu\lambda\underline{p}_s$.

(a) \underline{p}_c increases in μ . We have $\frac{\partial \underline{p}_c}{\partial \mu} > 0$ because $\frac{\partial \gamma}{\partial \mu} < 0$.

(b) \underline{p}_s decreases in μ . The derivative is $\frac{\partial \underline{p}_s}{\partial \mu} = -\frac{v(N-1)(1-\mu)^{N-2}}{(1+\gamma N)^2} < 0$.

(c) π increases in μ for all $\mu < \check{\mu}$ and decreases for $\mu \geq \check{\mu}$. The derivative is $\frac{\partial \pi}{\partial \mu} = \frac{\lambda \underline{p}_s}{(1-\mu)(1+\gamma N)} [1 + (1-\mu)\gamma N - \mu N]$. The RHS decreases in μ so $\frac{\partial \pi}{\partial \mu}$ is positive for $\mu < \check{\mu}$ and negative for $\mu > \check{\mu}$ where $\check{\mu}$ solves $(1-\lambda)(1-\check{\mu}N) + \lambda(1-\check{\mu})^N N = 0$, $\check{\mu} \in (0, 1)$. \square

Proof of Proposition 3. The comparative statics with respect to μ_s .

(i) I first cover the case when ϕ_n has an interior solution. I show in (a) that this holds for small μ_s : iff $\mu_s < \bar{\mu}_s$.

(a) \bar{p}_c and \underline{p}_c increase in μ_s . If ϕ_n has an interior solution, \bar{p}_c satisfies equation (13). Differentiating equation (13) with respect to μ_s yields $\frac{\partial \bar{p}_c}{\partial \mu_s} = \bar{p}_c(1-K)^{-1} \frac{\partial K}{\partial \gamma} \frac{\partial \gamma}{\partial \mu_s}$, where $\frac{\partial K}{\partial \gamma} = -\int_0^1 (1 + \gamma N y^{N-1})^{-2} N y^{N-1} dy < 0$ and $\frac{\partial \gamma}{\partial \mu_s} = -\frac{N\gamma}{1-\mu_s} < 0$ so $\frac{\partial \bar{p}_c}{\partial \mu_s} > 0$. As $\lim_{\mu_s \rightarrow 1} \bar{p}_c = +\infty$, ϕ_n is interior iff $\mu_s < \bar{\mu}_s$.

Also \underline{p}_c increases in μ_s . Inspecting equation (12) reveals that \underline{p}_c increases in \bar{p}_c and decreases in γ so $\frac{\partial \underline{p}_c}{\partial \mu_s} > 0$ because $\frac{\partial \bar{p}_c}{\partial \mu_s} > 0$ and $\frac{\partial \gamma}{\partial \mu_s} < 0$.

(b) \underline{p}_s increases in μ_s . The lowest shoppers' price is $\underline{p}_s = (1-\mu_s)^{N-1} \bar{p}_c (1+\gamma N)^{-1}$. Taking the derivative, using $\frac{\partial \bar{p}_c}{\partial \mu_s}$ from (a), collecting terms, and dividing the resulting equation with $\frac{(1-\mu_s)^{N-2} \bar{p}_c \gamma N}{(1+\gamma N)^2 (1-K)}$ gives

$$\frac{\partial \underline{p}_s}{\partial \mu_s} \propto \int_0^1 \frac{y^{N-1} [\gamma N^2 (1 - y^{N-1}) + (1 + \gamma N)(1 + \gamma N y^{N-1})]}{(1 + \gamma N y^{N-1})^2} dy > 0.$$

(c) π increases in μ_s . Since the profit is $\pi = \frac{1-\lambda}{N} \bar{p}_c + \mu_s \lambda \underline{p}_s$, $\frac{\partial \pi}{\partial \mu_s} > 0$.

(ii) I now cover the case when ϕ_n has a boundary solution, i.e., $\mu_s \geq \bar{\mu}_s$. If $\phi_n = v$, $\bar{p}_c = v$, $\underline{p}_c = v(1 + \gamma N)^{-1}$, $\underline{p}_s = v(1 - \mu_s)^{N-1}(1 + \gamma N)^{-1}$, and $\pi = \frac{1-\lambda}{N}v + \lambda\mu_s\underline{p}_s$.

(a) \underline{p}_c increases in μ_s . We have $\frac{\partial \underline{p}_c}{\partial \mu_s} > 0$ because $\frac{\partial \gamma}{\partial \mu_s} < 0$.

(b) \underline{p}_s increases in μ_s for all $\mu_s < \check{\mu}_{s1}$ and decreases for all $\mu_s \geq \check{\mu}_{s1}$. The derivative $\frac{\partial \underline{p}_s}{\partial \mu_s} = \frac{v(1-\mu_s)^{N-2}}{(1+\gamma N)^2} (1 + \gamma N - N)$, is positive for all $\mu_s < \check{\mu}_{s1}$ where $\check{\mu}_{s1} := 1 - \left[\frac{(N-1)(1-\lambda)(1-\mu_n)}{N\lambda} \right]^{\frac{1}{N}}$. We have $\check{\mu}_{s1} < 1$, but not necessarily $\check{\mu}_{s1} > 0$.

- (c) π increases in μ_s for $\mu_s < \check{\mu}_{s2}$ and decreases for $\mu_s \geq \check{\mu}_{s2}$. The derivative is $\frac{\partial \pi}{\partial \mu_s} = \frac{\lambda \underline{p}_s}{(1-\mu_s)(1+\gamma N)}(1 + \gamma N - \mu_s N)$, which is positive for all $\mu_s < \check{\mu}_{s2}$ where $\check{\mu}_{s2}$ solves $\check{\mu}_{s2} - \frac{\lambda(1-\check{\mu}_{s2})^N}{(1-\lambda)(1-\mu_n)} = N^{-1}$ with $\check{\mu}_{s2} \in (0, 1)$ and $\check{\mu}_{s1} < \check{\mu}_{s2}$.

The comparative statics with respect to μ_n .

- (i) I first cover the case when ϕ_n has an interior solution. I show in (a) that this holds for small μ_n : iff $\mu_n < \bar{\mu}_n$.

- (a) \bar{p}_c increases and \underline{p}_c decreases in μ_n . Differentiating equation (13) yields

$$\frac{\partial \bar{p}_c}{\partial \mu_n} = \frac{\partial \phi_n}{\partial \mu_n} = \frac{\bar{p}_c}{(1-\mu_n)(1-K)} \int_0^1 \left(\frac{\gamma N y^{N-1}}{1 + \gamma N y^{N-1}} \right)^2 dy > 0. \quad (15)$$

As $\lim_{\mu_n \rightarrow 1} \bar{p}_c = +\infty$, ϕ_n has an interior solution iff $\mu_n < \bar{\mu}_n$.

For \underline{p}_c , I take the derivative of (12), plug in $\frac{\partial \bar{p}_c}{\partial \mu_n}$ and collect terms, to get

$$\frac{\partial \underline{p}_c}{\partial \mu_n} = -\frac{\bar{p}_c \gamma^2 N^2}{(1 + \gamma N)^2 (1 - \mu_n)(1 - K)} \int_0^1 \frac{y^{N-1}(1 - y^{N-1})}{(1 + \gamma N y^{N-1})^2} dy < 0. \quad (16)$$

- (b) \underline{p}_s decreases in μ_n . As $\underline{p}_s = (1 - \mu_s)^{N-1} \underline{p}_c$ and $\frac{\partial \underline{p}_c}{\partial \mu_n} < 0$, also $\frac{\partial \underline{p}_s}{\partial \mu_n} < 0$.
- (c) A sufficient condition for π to increase in μ_n is that $\mu_n \leq \check{\mu}_{n1}$. The expected equilibrium profit is $\pi = \frac{1-\lambda}{N} \bar{p}_c + \mu_s \lambda \underline{p}_s$. Using equations (15) and (16) collecting terms, and dividing by $\frac{\bar{p}_c(1-\lambda)\gamma^2 N}{(1-\mu_n)(1-\mu_s)(1-K)(1+\gamma N)^2}$, I get

$$\frac{\partial \pi}{\partial \mu_n} \propto \int_0^1 \frac{y^{N-1} [(1 - \mu_s)(1 + \gamma N)^2 y^{N-1} - \mu_s(1 - \mu_n)\gamma N(1 - y^{N-1})]}{(1 + \gamma N y^{N-1})^2} dy.$$

Denote the integrand by Φ . I want to show that $\int_0^1 \Phi dy > 0$. Note that

$$\frac{\partial \Phi}{\partial \mu_n} = \frac{2y^{2(N-1)}N(1 - y^{N-1})\gamma}{(1 + \gamma N y^{N-1})^3(1 - \mu_n)} [(1 - \mu_s)(1 + \gamma N) + \mu_s(1 - \mu_n)\gamma N] > 0,$$

so that $\frac{\partial}{\partial \mu_n} \int_0^1 \Phi dy > 0$. Instead of showing that $\int_0^1 \Phi dy > 0$, I derive a sufficient condition for $\int_0^1 \Phi|_{\mu_n=0} dy > 0$. I approximate down

$$\Phi|_{\mu_n=0} = \frac{(1 - \mu_s)(1 + \gamma N)^2 y^{2N-2}}{(1 + \gamma N y^{N-1})^2} + \frac{\mu_s \gamma N y^{2N-2}}{(1 + \gamma N y^{N-1})^2} - \frac{\mu_s \gamma N y^{N-1}}{(1 + \gamma N y^{N-1})^2}.$$

In $\Phi|_{\mu_n=0}$, I substitute y^{N-1} by one in the denominators of the two first

fractions, and y^{N-1} by y^{N-2} in the numerator of the last fraction. This gives

$$W := (1 - \mu_s)y^{2N-2} + \frac{\mu_s\gamma Ny^{2N-2}}{(1 + \gamma N)^2} - \frac{\mu_s\gamma Ny^{N-2}}{(1 + \gamma Ny^{N-1})^2}.$$

I derive a sufficient condition for $\int_0^1 W dy > 0$. Direct evaluation gives

$$\int_0^1 W dy = \frac{1 - \mu_s}{2N - 1} - \frac{\mu_s\gamma N^2[1 + \gamma(2N - 1)]}{(1 + \gamma N)^2(2N - 1)(N - 1)}.$$

The RHS decreases in γ , $RHS|_{\gamma=0} > 0$ and $\lim_{\gamma \rightarrow \infty} RHS = \frac{(N-1)(1-\mu_s)-(2N-1)\mu_s}{(N-1)(2N-1)}$. If μ_s is small, $\int_0^1 W dy > 0$ holds for all γ and if μ_s is large, for $\gamma \leq \check{\gamma}_1$ where $\check{\gamma}_1$ solves $\frac{\partial}{\partial \gamma} \int_0^1 W dy = 0$. I rewrite $\gamma \leq \check{\gamma}_1$ as $\mu_n \leq \check{\mu}_{n1}$ for large μ_s and let $\check{\mu}_{n1} = 1$ for small μ_s .

(ii) I now cover the case when ϕ_n has a boundary solution, i.e., $\mu_n > \bar{\mu}_n$. If $\phi_n = v$, then $\bar{p}_c = v$, $\underline{p}_c = v(1 + \gamma N)^{-1}$, $\underline{p}_s = (1 - \mu_s)^{N-1}\underline{p}_c$ and $\pi = \frac{1-\lambda}{N}v + \lambda\mu_s\underline{p}_s$.

(a)-(c) \underline{p}_c , \underline{p}_s and π decrease in μ_n . All follow directly from $\frac{\partial \gamma}{\partial \mu_n} = \frac{\gamma}{1-\mu_n} > 0$. \square

Proof of Proposition 4. The comparative statics with respect to N .

(i) I first cover the case when ϕ_n has an interior solution. I show in (a) that this holds for small N : iff $N < \bar{N}$.

(a) \bar{p}_c increases in N and a sufficient condition for \underline{p}_c to increase is $N \geq \check{N}_1$. Directly taking the derivative of equation (13) and collecting terms gives

$$\frac{\partial \bar{p}_c}{\partial N} = -\frac{\bar{p}_c}{1 - \int_0^1 (1 + N\gamma y^{N-1})^{-1} dy} \int_0^1 \frac{\gamma y^{N-1}[N \ln(1 - \mu_s) + 1 + N \ln y]}{(1 + \gamma N y^{N-1})^2} dy. \quad (17)$$

I use a result from Stahl (1989) to show that $\frac{\partial \bar{p}_c}{\partial N} > 0$. Let the highest price in the unit-demand version of Stahl (1989) be \bar{p}_{stahl} . An interior \bar{p}_{stahl} can be written as $\bar{p}_{stahl} = \alpha \left[1 - \int_0^1 (1 + \hat{\gamma} N y^{N-1})^{-1} dy\right]^{-1}$, where $\hat{\gamma} := \frac{\lambda}{1-\lambda}$ (Janssen et al., 2005). Stahl (1989) shows that

$$\frac{\partial \bar{p}_{stahl}}{\partial N} = -\frac{\bar{p}_{stahl}}{1 - \int_0^1 (1 + N\hat{\gamma} y^{N-1})^{-1} dy} \int_0^1 \frac{\hat{\gamma} y^{N-1}(1 + N \ln y)}{(1 + \hat{\gamma} N y^{N-1})^2} dy > 0. \quad (18)$$

Since $\frac{\partial \bar{p}_{stahl}}{\partial N} > 0$, the last integral in (18) is negative. Both $\hat{\gamma}$ and γ can take any value in $[0, 1]$ so to show that the integral in equation (17) is negative, it suffices to show that the integrand in (17) is below that in (18) for all $y \in [0, 1]$. This holds because $N \ln(1 - \mu_s) \leq 0$ for all $\mu_s \in [0, 1]$.

Recall that $\underline{p}_c = \bar{p}_c(1 + \gamma N)^{-1}$. The derivative

$$\frac{\partial \underline{p}_c}{\partial N} = \frac{\partial \bar{p}_c}{\partial N}(1 + \gamma N)^{-1} + \frac{\gamma \bar{p}_c}{(1 + \gamma N)^2}[-N \ln(1 - \mu_s) - 1],$$

is positive for sure if the last term is positive: if $N \geq (-\ln(1 - \mu_s))^{-1} =: \check{N}_1$.

(b) \underline{p}_s decreases and increases in N for nonempty open sets of parameter values.

Recall that $\underline{p}_s = (1 - \mu_s)^{N-1} \frac{\bar{p}_c}{1 + \gamma N}$, so that

$$\frac{\partial \underline{p}_s}{\partial N} = \frac{(1 - \mu_s)^{N-1}}{(1 + \gamma N)^2} \left\{ \bar{p}_c[\ln(1 - \mu_s) - \gamma] + (1 + \gamma N) \frac{\partial \bar{p}_c}{\partial N} \right\}. \quad (19)$$

Let $\hat{X} := (N = 2, \gamma = \frac{1}{2}, \mu_s = \mu_n = \hat{\mu})$ where $\hat{\mu}$ solves $\ln(1 - \hat{\mu}) = -\frac{1}{2}$.

Evaluating (19) at \hat{X} and collecting terms gives

$$\frac{\partial \underline{p}_s}{\partial N} \Big|_{\hat{X}} \propto - \int_0^1 \frac{y(1+y) + 2y \ln y}{(1+y)^2} dy = - \left(1 + \ln 2 - \frac{\pi^2}{6} \right) < 0.$$

I now show that parameter values exist for which $\frac{\partial \underline{p}_s}{\partial N} > 0$. Let $\tilde{X} := (N = 2, \gamma = \frac{1}{2}, \mu_s = \mu_n = \tilde{\mu})$ where $\tilde{\mu}$ solves $\ln(1 - \tilde{\mu}) = -2$. Evaluating (19) at \tilde{X} and collecting terms gives

$$\frac{\partial \underline{p}_s}{\partial N} \Big|_{\tilde{X}} \propto - \int_0^1 \frac{y(-1+5y) + 4y \ln y}{(1+y)^2} dy = - \left(8 - 7 \ln 2 - \frac{\pi^2}{3} \right) > 0.$$

Neither inequality is tight and everything is continuous around \hat{X} and \tilde{X} , so both comparative statics hold for nonempty open sets of parameter values.

(c) Sufficient conditions for π to increase in N are $N > 4$, $\mu_s \geq \check{\mu}_{s3}$, and $\mu_n \in (\check{\mu}_{n2}, \check{\mu}_{n3})$. I proceed in three steps (c.i)-(c.iii). I first rewrite the inequality $\pi_{N+1} > \pi_N$. I then show that sufficient conditions for the new inequality's LHS to exceed $\frac{5}{4}$ are $\mu_s \geq \check{\mu}_{s3}$ and $\mu_n \leq \check{\mu}_{n3}$, and for the RHS to be below $\frac{5}{4}$ are $N > 4$ and $\mu_n \geq \check{\mu}_{n2}$.

(c.i) $\pi_{N+1} > \pi_N$ can be rewritten as inequality (20). Using the profit equations gives that $\pi_{N+1} > \pi_N$ is equivalent to

$$\begin{aligned} & \left\{ \frac{[1 + \gamma_{N+1}(N+1)](1 - \mu_s) + \mu_s \gamma_{N+1}(N+1)(1 - \mu_n)}{(N+1)[1 + \gamma_{N+1}(N+1)]} \right\} \times \bar{p}_c|_{N+1} \\ & \geq \left[\frac{(1 + \gamma_N N)(1 - \mu_s) + \mu_s \gamma_N N(1 - \mu_n)}{N(1 + \gamma_N N)} \right] \times \bar{p}_c|_N \end{aligned}$$

where $\bar{p}_c|_{N+1}$ is the highest common price when the number of firms is $N + 1$ and $\gamma_{N+1} = \frac{\lambda}{1-\lambda} \frac{(1-\mu_s)^{N+1}}{1-\mu_n}$. Since $\gamma_{N+1} = \frac{\lambda}{1-\lambda} \frac{(1-\mu_s)^{N+1}}{1-\mu_n} = (1 - \mu_s)\gamma_N$, the inequality can be rewritten as

$$\frac{\bar{p}_c|_{N+1}}{\bar{p}_c|_N} \geq \left[\frac{(1 + \gamma_N N)(1 - \mu_s) + \mu_s \gamma_N N(1 - \mu_n)}{N(1 + \gamma_N N)} \right] \quad (20)$$

$$\times \left\{ \frac{(N + 1)[1 + (1 - \mu_s)\gamma_N(N + 1)]}{(1 - \mu_s)[1 + (1 - \mu_s)\gamma_N(N + 1) + \mu_s \gamma_N(N + 1)(1 - \mu_n)]} \right\} =: RHS_{(20)}.$$

(c.ii) A sufficient condition for $\frac{\bar{p}_c|_{N+1}}{\bar{p}_c|_N} \geq \frac{5}{4}$, or $4\bar{p}_c|_{N+1} \geq 5\bar{p}_c|_N$, is $\mu_s \geq \check{\mu}_{s3}$ and $\mu_n \leq \check{\mu}_{n3}$. I use equation (13) to rewrite the inequality $4\bar{p}_c|_{N+1} \geq 5\bar{p}_c|_N$ as

$$\int_0^1 \frac{5}{1 + \gamma_N(N + 1)(1 - \mu_s)y^N} dy \geq \int_0^1 \frac{5 + \gamma_N N y^{N-1}}{1 + \gamma_N N y^{N-1}} dy. \quad (21)$$

A sufficient condition for (21) to hold is that the integrand on the LHS exceeds that on the RHS, which holds if $4N \geq (5 + \gamma_N N)(N + 1)(1 - \mu_s)y$ for all $y \in [0, 1]$. The latter is the hardest to satisfy for $y = 1$ and holds for $y = 1$ if

$$4N \geq \left[5 + \frac{\lambda}{1-\lambda} \frac{(1 - \mu_s)^N}{1 - \mu_n} N \right] (N + 1)(1 - \mu_s), \quad (22)$$

which gives a joint condition on μ_n and μ_s .²⁶ Since the RHS in (22) decreases in μ_s and increases in μ_n , the inequality holds if $\mu_s \geq \check{\mu}_{s3}$ and $\mu_n \leq \check{\mu}_{n3}$.

(c.iii) Sufficient conditions for $\frac{5}{4}$ to weakly exceed the RHS of equation (20) are $N > 4$ and $\mu_n \geq \check{\mu}_{n2}$. The inequality $\frac{5}{4} \geq RHS_{(20)}$ can be rewritten as

$$(N - 4)(1 - \mu_s)[1 + \gamma N + \gamma^2 N(N + 1)(1 - \mu_n \mu_s)]$$

$$+ \gamma N(N + 1)(1 - \mu_n \mu_s)(1 - 5\mu_s) + 4\gamma(N + 1)(1 - \mu_s)[(N + 1)\mu_s - 1] \geq 0.$$

The first term of this inequality is weakly positive for all $N \geq 4$ so we need to consider the last two terms. If $\mu_s \leq \frac{1}{5}$, the terms' sum is positive because $N \geq 4$, $(1 - \mu_n \mu_s) \geq (1 - \mu_s)$ and $1 - 5\mu_s \geq (N + 1)\mu_s - 1$. If $\mu_s > \frac{1}{5}$, the terms' sum can be negative. Sufficient conditions for the entire inequality to still hold for all μ_s are $N > 4$ and $\mu_n \geq \check{\mu}_{n2}$. For large N , $\check{\mu}_{n2}$ is zero.

(d) $\bar{p}_c - \underline{p}_s$ increases in N . Since $\bar{p}_c - \underline{p}_s = \bar{p}_c [1 - (1 - \mu_s)^{N-1}(1 + \gamma N)^{-1}]$, we

²⁶If the inequality does not hold for all $y \in [0, 1]$, inequality (21) can still hold and, everything else constant, is easier to satisfy for large μ_s .

have

$$\begin{aligned} \frac{\partial(\bar{p}_c - \underline{p}_s)}{\partial N} &= \bar{p}'_c [1 - (1 - \mu_s)^{N-1}(1 + \gamma N)^{-1}] \\ &+ \bar{p}_c(1 - \mu_s)^{N-1}(1 + \gamma N)^{-2} [-(1 + \gamma N) \ln(1 - \mu_s) + \gamma(N \ln(1 - \mu_s) + 1)]. \end{aligned}$$

Using \bar{p}'_c from (a) above gives

$$\begin{aligned} \frac{\partial(\bar{p}_c - \underline{p}_s)}{\partial N} &\propto - [(1 + \gamma N) - (1 - \mu_s)^{N-1}] \int_0^1 \frac{y^{N-1} [N \ln(1 - \mu_s) + 1 + N \ln y]}{(1 + \gamma N y^{N-1})^2} dy \\ &+ (1 - \mu_s)^{N-1}(1 + \gamma N)^{-1} [\gamma - \ln(1 - \mu_s)] \int_0^1 \frac{N y^{N-1}}{1 + N \gamma y^{N-1}} dy > 0 \end{aligned}$$

because, we know from (a) above, the first integral is negative.

(ii) I now cover the case when ϕ_n has a boundary solution, i.e., $N \geq \bar{N}$. If $\phi_n = v$, then $\bar{p}_c = v$, $\underline{p}_c = v(1 + \gamma N)^{-1}$, $\underline{p}_s = (1 - \mu_s)^{N-1} \underline{p}_c$, and $\pi = \frac{1-\lambda}{N}v + \lambda \mu_s \underline{p}_s$.

- (a) \underline{p}_c increases in N for all $N > \check{N}_1$ and decreases for all $N \leq \check{N}_1$. The derivative $\frac{\partial \underline{p}_c}{\partial N} = \frac{v\gamma}{(1+\gamma N)^2} [-N \ln(1 - \mu_s) - 1]$ is positive iff $N > (-\ln(1 - \mu_s))^{-1} = \check{N}_1$.
- (b) \underline{p}_s decreases in N . The derivative is $\frac{\partial \underline{p}_s}{\partial N} = \frac{\underline{p}_c}{1+\gamma N} (1 - \mu_s)^{N-1} [\ln(1 - \mu_s) - \gamma] < 0$.
- (c) π decreases in N . We have $\frac{\partial \pi}{\partial N} < 0$ because $\frac{\partial \underline{p}_s}{\partial N} < 0$.
- (d) $\bar{p}_c - \underline{p}_s$ increases in N . Follows from $\bar{p}_c = v$ and $\frac{\partial \underline{p}_s}{\partial N} < 0$. \square

A.3 Example screenshots

Figures 5 and 6 are screenshots of small parts of Amazon's privacy policy. Figure 5 highlights some aspects of a person's behaviour that the firm records and that may be indicative of how much a consumer looks around before buying. In particular, the information includes not only how long the person spends on different pages and what he searches for, but also page interaction data such as scrolling, clicks and even mouse-overs. The collected information includes which website the person arrives from before (and goes to after) visiting Amazon's site.

Figure 6 shows that other firms can access information about a person's interactions with Amazon's website using third-party cookies. It also shows that a person's behaviour can be followed for a year or longer because cookies are stored on a person's device for 13 months (after a single approval of cookies).

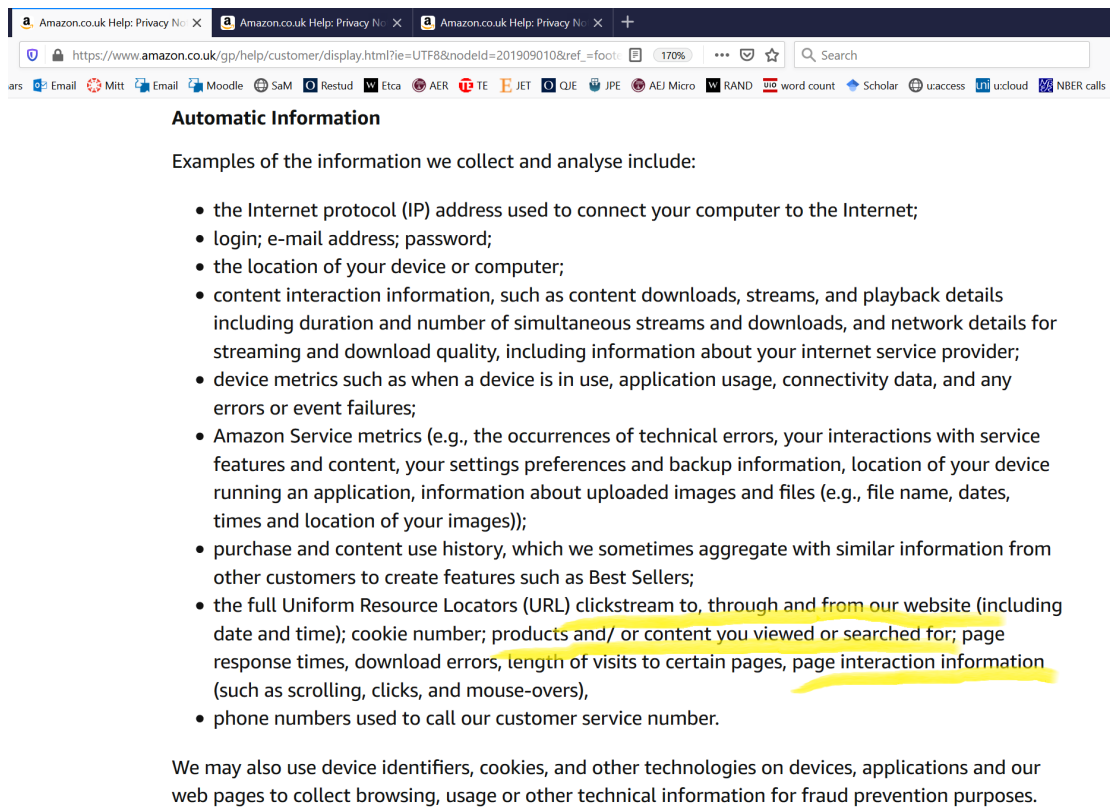


Figure 5: An extract from Amazon's Privacy Policy (screenshot from 28.09.20).

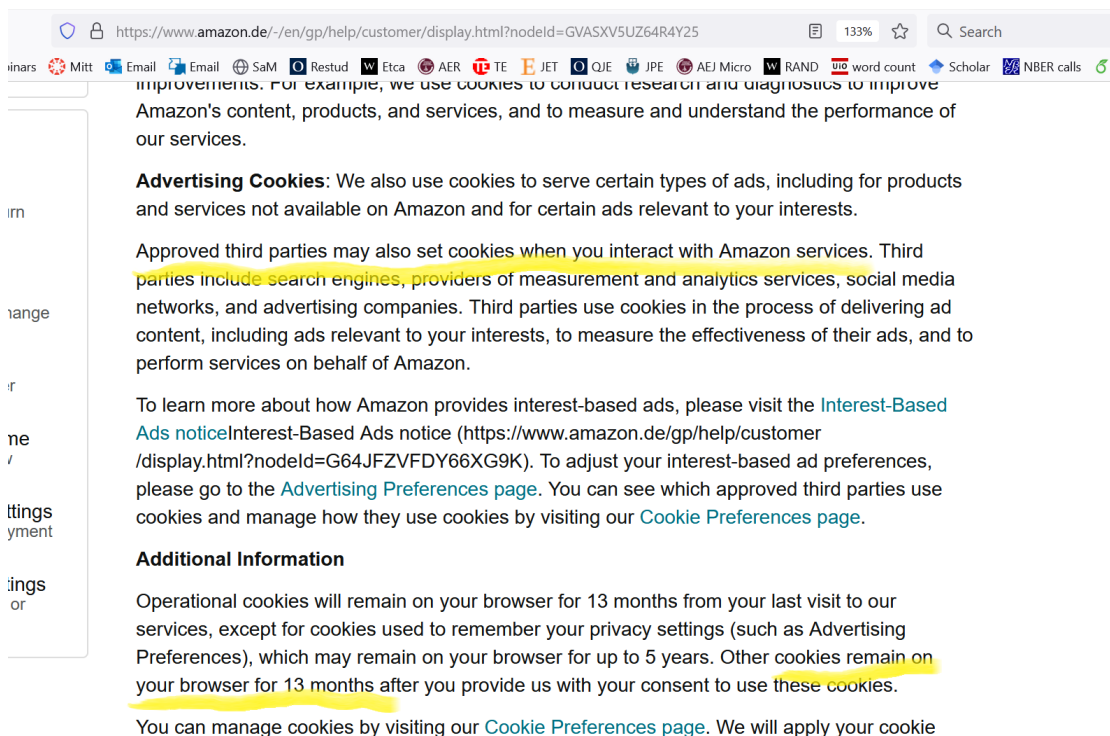


Figure 6: An extract from Amazon's Privacy Policy (screenshot from 16.02.22).

References

- ACQUISTI, A. AND H. VARIAN (2005): “Conditioning Prices on Purchase History,” *Marketing Science*, 24, 367–381.
- AMAZON (2021): “Privacy Policy – Amazon,” <https://www.amazon.com/gp/help/customer/display.html?nodeId=GX7NJQ4ZB8MHFRNJ>, [Online; accessed 28.09.20, 04.04.21, 16.02.22].
- ARMSTRONG, M. AND J. VICKERS (2019): “Discriminating against Captive Customers,” *American Economic Review: Insights*, 1, 257–72.
- ARMSTRONG, M. AND J. ZHOU (2016): “Search Deterrence,” *Review of Economic Studies*, 83, 26–57.
- ATAYEV, A. (2020): “Search Cost Based Price Discrimination,” Working Paper.
- BELLEFLAMME, P., W. LAM, AND W. VERGOTE (2020): “Competitive imperfect price discrimination and market power,” *Marketing Science*, 39, 996–1015.
- BERGEMANN, D., B. BROOKS, AND S. MORRIS (2021): “Search, Information, and Prices,” *Journal of Political Economy*, 129, 2275–2319.
- BÖHEIM, R., F. HACKL, AND M. HÖLZL-LEITNER (2021): “The Impact of Price Adjustment Costs on Price Dispersion in E-commerce,” *International Journal of Industrial Organization*, 77, 102743.
- BOURREAU, M. AND A. DE STREEL (2018): “The Regulation of Personalised Pricing in the Digital Era,” OECD Meeting Note.
- BRAGHERI, L. (2019): “Targeted Advertising and Price Discrimination in Intermediated Online Markets,” Working Paper.
- CHEN, Y. (1997): “Paying Customers to Switch,” *Journal of Economics & Management Strategy*, 6, 877–897.
- CONITZER, V., C. TAYLOR, AND L. WAGMAN (2012): “Hide and Seek: Costly Consumer Privacy in a Market with Repeat Purchases,” *Marketing Science*, 31, 277–292.
- DE LOS SANTOS, B., A. HORTAÇSU, AND M. R. WILDENBEEST (2012): “Testing models of consumer search using data on web browsing and purchasing behavior,” *American Economic Review*, 102, 2955–80.
- ENNIS, S. AND W. LAM (2021): “Personalised Pricing and Disclosure,” BEIS Research Paper Number 2021/008.
- EUROSTAT (2021): “Individuals - Mobile Internet Access,” <http://appsso.eurostat.ec.europa.eu/nui/show.do>, [Online; accessed: 16.09.21].
- FABRA, N. AND M. REGUANT (2020): “A Model of Search with Price Discrimination,” *European Economic Review*, 129.

- FACEBOOK (2021): “Privacy Policy – Facebook,” <https://www.facebook.com/about/privacy/>, [Online; accessed 28.09.20, 04.04.21].
- FCA (2019): “General Insurance Pricing Practices,” Financial Conduct Authority Market Study MS18/1.2.
- FUDENBERG, D. AND J. TIROLE (2000): “Customer Poaching and Brand Switching,” *RAND Journal of Economics*, 634–657.
- GOOGLE (2021): “Privacy Policy – Google,” <https://policies.google.com/privacy?hl=en-US#whycollect>, [Online; accessed 28.09.20].
- GREWAL, D., R. JANAKIRAMAN, K. KALYANAM, P. KANNAN, B. RATCHFORD, R. SONG, AND S. TOLERICO (2010): “Strategic Online and Offline Retail Pricing: A Review and Research Agenda,” *Journal of Interactive Marketing*, 24, 138–154.
- HANNAK, A., G. SOELLER, D. LAZER, A. MISLOVE, AND C. WILSON (2014): “Measuring Price Discrimination and Steering on E-Commerce Web Sites,” in *Proceedings of the 2014 Conference on Internet Measurement*, 305–318.
- HART, O. D. AND J. TIROLE (1988): “Contract Renegotiation and Coasian Dynamics,” *Review of Economic Studies*, 55, 509–540.
- HAYNES, M. AND S. THOMPSON (2008): “Price, Price Dispersion and Number of Sellers at a Low Entry Cost Shopbot,” *International Journal of Industrial Organization*, 26, 459–472.
- HIDIR, S. AND N. VELLODI (2021): “Privacy, Personalization, and Price Discrimination,” *Journal of the European Economic Association*, 1342–1363.
- HONKA, E. (2014): “Quantifying search and switching costs in the US auto insurance industry,” *The RAND Journal of Economics*, 45, 847–884.
- ICHIHASHI, S. (2020): “Online Privacy and Information Disclosure by Consumers,” *American Economic Review*, 110, 569–95.
- IPSOS, LONDON ECONOMICS, AND DELOITTE CONSORTIUM (2018): “Consumer Market Study on Online Market Segmentation Through Personalised Pricing/Offers in the European Union,” European Commission Report.
- JANSSEN, M., J. L. MORAGA-GONZÁLEZ, AND M. WILDENBEEST (2005): “Truly Costly Sequential Search and Oligopolistic Pricing,” *International Journal of Industrial Organization*, 23, 451–466.
- LOERTSCHER, S. AND L. MARX (2020): “Digital monopolies: Privacy protection or price regulation?” *International Journal of Industrial Organization*, 71, 102623.
- MONTES, R., W. SAND-ZANTMAN, AND T. VALLETTI (2019): “The Value of Personal Information in Online Markets with Endogenous Privacy,” *Management Science*, 65, 1342–1362.

- NELSON, R., R. COHEN, AND F. RASMUSSEN (2007): “An Analysis of Pricing Strategy and Price Dispersion on the Internet,” *Eastern Economic Journal*, 33, 95–110.
- PREUSS, M. (2021): “Search, Learning and Tracking,” *RAND Journal of Economics*, forthcoming.
- STAHL, D. (1989): “Oligopolistic Pricing with Sequential Consumer Search,” *American Economic Review*, 700–712.
- (1996): “Oligopolistic Pricing with Heterogeneous Consumer Search,” *International Journal of Industrial Organization*, 14, 243–268.
- TAYLOR, C. R. (2003): “Supplier Surfing: Competition and Consumer Behavior in Subscription Markets,” *RAND Journal of Economics*, 223–246.
- VILLAS-BOAS, J. M. (1999): “Dynamic competition with customer recognition,” *RAND Journal of Economics*, 604–631.
- WANG, W., F. LI, AND Y. ZHANG (2021): “Price Discount and Price Dispersion in Online Market: Do More Firms Still Lead to More Competition?” *Journal of Theoretical and Applied Electronic Commerce Research*, 16, 140–154.