Sources of limited consideration and market power in e-commerce^{*}

Michael R. Sullivan

University of Western Ontario

October 10, 2024

Click here for the online appendix.

Abstract

This article develops techniques for the empirical analysis of repeated sequential search over unordered alternatives using data on consumer search processes. I use these techniques to assess why consumers conduct little search in e-commerce and often pay significantly above the minimum available price for a product. Search costs could explain these facts, as could pre-search seller differentiation: consumers with low search costs may not visit stores they dislike based on information known before search. I find that seller differentiation is primarily responsible for limited consideration and market power.

^{*}Email: msulli65@uwo.ca. I thank Steven Berry, Philip Haile, Katja Seim, Mitsuru Igami, and Soheil Ghili as well as seminar participants at Yale, the Young Economist Symposium, the 5th Doctoral Workshop on the Economics of Digitization, the 2024 Econometric Society North American Winter Meetings, and the 2024 International Industrial Organization Conference for helpful feedback. This article draws on research supported by the Social Sciences and Humanities Research Council. Martin Bustos provided capable research assistance.

1 Introduction

This article develops techniques for the empirical analysis of directed sequential search. These techniques apply to the setting in which a consumer sequentially searches unordered alternatives and the search process (i.e., the identities and ordering of searched alternatives) is observed by the researcher. The starting point of the analysis is the Weitzman (1979) optimal sequential search strategy. I show that this strategy implies a one-to-one mapping between search effort outcomes and chains of inequalities relating consumer utilities for an arbitrary number of alternatives. This mapping facilitates model analysis even without parametric restrictions on search costs. However, the mapping is especially useful under a particular search cost distribution proposed in the article and inspired by the approach of Moraga-González et al. (2023). The combination of this distribution and the mapping between search effort outcomes and utility inequalities yields closed-form expressions for joint search and purchase outcomes. This is the case even when the distribution of search costs varies across retailers. One advantage of the model is that it permits analysis of the relative roles of awareness and quality differentials in explaining retailers' market shares.

The article's techniques can be used to quantify how search frictions, vertical differentiation, and horizontal differentiation determine market dominance and market power in e-commerce markets. The primary applied question I address using these techniques is a fundamental one in e-commerce: what drives limited search and retailer market power in online markets for minimally differentiated products? If internet search were costless and both sellers and their product offerings were undifferentiated, consumers would compare prices across all retailers and purchase from the lowest-price seller. In reality, however, consumers conduct often little search and buy from higher-price sellers of undifferentiated products.¹ This could reflect that search fric-

¹See, e.g., Clay et al. (2001), Clemons et al. (2002), Moraga-González and Wildenbeest

tions remain significant online — consumers may buy at high prices to avoid further search. Much of the empirical online search literature emphasizes this explanation.² Seller differentiation can also explain limited consideration. Even when the product that arrives on a consumer's doorstep does not vary across retailers, a consumer may differentially value retailers due to vertical differences in shipping efficiency or customer service. Retailers may also be horizontally differentiated by their user interfaces and marketing strategies. Additionally, consumers may prefer to buy from stores that they have previously patronized due to habit formation, store loyalty, or switching costs. If the consumer knows before searching that they are unlikely to buy from a seller, then the consumer may not visit the seller even when search costs are negligible.

I empirically investigate sources of limited search and market power in US contact lens e-commerce. This setting is attractive for the study of acrossretailer search because consumers require brand-specific prescriptions to buy lenses, which allows me to credibly assume that search occurs across stores and not across products. With that said, the article's framework is readily adapted to the analysis of search across products when the researcher possesses data on such search. The article's methodological innovations facilitate estimation and analysis of a model of sequential search for contact lenses repeated over time. The availability of data on both store visits and on purchases permits a quantification of the distinct roles of awareness and quality differences in determining retailers' sales. One challenge in identifying parameters affecting purchase utility is price endogeneity, which owes to the dependence of unobserved retailer quality and prices. My solution to this problem exploits within-retailer, across-brand variation in relative prices and relative market shares. Under this solution, the extent to which a retailer

^{(2008),} Koulayev (2014), and Jolivet and Turon (2019).

 $^{^{2}}$ See Hortaçsu and Syverson (2004), Hong and Shum (2006), Moraga-González and Wildenbeest (2008), and Jolivet and Turon (2019).

has a relatively low market share in sales of a brand that it sells for a relatively high price identifies price sensitivity. Separate identification of state dependence and unobserved heterogeneity follows from standard arguments concerning their distinct implications for choice dynamics.

Indirect-inference estimates of the model imply median search costs of under \$1.25 for all retailers. Removing various forms of seller differentiation from the model raises estimated search costs dramatically — eliminating factors that limit search requires search costs to play a larger role in justifying limited consideration. This finding suggests that flexible modelling of retailer differentiation is essential in reliably estimating search costs.

Both search frictions and store differentiation play a role in limiting search. Eliminating vertical differentiation—i.e., differences in mean consumer tastes for retailers—raises the mean number of store visits from 1.20 to 1.30 by inducing consumers who prefer the vertically superior retailer to consider other stores. Eliminating horizontal differentiation—i.e., cross-consumer dispersion in tastes for stores—similarly boosts search intensity by leading consummers to look beyond their favoured store. Although reducing search costs raises search intensity, it does not meaningfully affect the extent to which consumers pay above the minimum available price for contacts. I instead find that consumers pay above the minimum available price largely because they value the superior quality of higher-price retailers, which could reflect superior shipping times, customer service, or return policies. I additionally assess sources of markups in contact lens e-commerce. The results suggest that seller differentiation shapes equilibrium markups whereas search costs do not. For one popular brand, eliminating horizontal differentiation reduces markups by 55% on average. Additionally, eliminating the upscale retailer's vertical advantage reduces its markups by 20%. Prices at rival retailers rise absent vertical differentiation, thus reducing price dispersion. Results for

other brands are similar. Together, the results suggest that retailer differentiation is responsible for market power and price dispersion in e-commerce.

The model also permits an analysis of the roles of awareness and quality advantages in explaining a firm's dominance in a market. The largest contact lens retailer, 1-800 Contacts, outsold its rivals at higher prices. Notably, 1-800 Contacts was also known for its superior service quality and extensive advertising, which is reflected in a higher estimate of quality and a lower estimate of search costs for 1-800 than for its rivals. I show that 1-800's quality advantage rather than its awareness advantage underlies its market dominance: equalizing the quality of 1-800 Contacts and Vision Direct, its main rival, reduces the ratio of the former's sales to the latter's from 1.71 to 0.38, whereas equalizing these retailers' search cost distributions reduces the ratio only to 1.67.

Since 2007–2008, e-commerce has witnessed entry of many new sellers, including those that primarily sell within e-commerce platforms, and changes in the nature of advertising. The article's framework is well suited for the study of contemporary e-commerce markets. The model is easily adapted to the case of consumer search across products or third-party sellers on an e-commerce platform. Furthermore, the fact that article's characterization of the probabilities of joint search and purchase outcomes holds for any arbitrary number of retailers makes its methods useful in settings with many retailers. The model also permits heterogeneity in the magnitude of search costs across retailers and variance in consumer/retailer-specific search costs; this allows it to capture both (i) differences in the intensity of informative advertising across retailers and (ii) personalized advertisements, both features of contemporary e-commerce.

1.1 Related literature

The article's primary contribution is the development of techniques for estimating a sequential search model using data on consumer search processes, namely the establishment of (i) a one-to-one mapping between search outcomes and utility inequalities and (ii) a parametric specification that operationalizes this mapping. These techniques draw on Weitzman (1979) and Moraga-González et al. (2023). The article extends the analysis of Moraga-González et al. (2023), which provides expressions for probabilities of purchase outcomes, by providing expressions for probabilities of both search and purchase outcomes. These expressions are useful in drawing upon the identifying power of datasets that describe not only purchase decisions but also search processes. Absent the techniques in this article, complexities arise in the analysis of sequential consumer search with data on consumer search processes. Honka and Chintagunta (2017) provide a foundational study in the estimation of the sequential search model with data on search sequences. In a setting similar to my own, they pool together distinct sets of inequalities that characterize search effort outcomes and approximate the probabilities implied by these inequalities via simulation. I build upon their contribution by developing techniques that yield closed-form probabilities from comprehensive chains of inequalities characterizing both search and purchase.

I also develop econometric techniques for the analysis of search data with a panel dimension. Whereas recent studies have considered persistent unobserved heterogeneity (Morozov et al. 2021) and state dependence (Honka 2014) separately, my article considers both phenomena simultaneously, proposing solutions based on the panel econometrics literature to an endogeneity problem and an initial conditions problem that arise.

This article's applied contribution is its explanation of limited search and market power in e-commerce. Brynjolfsson and Smith (2000) studied price dispersion in early e-commerce, concluding that seller heterogeneity remained significant on the internet. Early articles in the empirical consumer search literature—namely Hong and Shum (2006), Hortaçsu and Syverson (2004), and Moraga-González and Wildenbeest (2008)—demonstrated that search frictions could explain price dispersion in homogeneous goods markets.³ Several recent studies account for other factors that limit search and generate market power both within and outside of e-commerce (e.g., Honka 2014, Morozov et al. 2021, Brown et al. 2023).

My article's methods are specialized to the setting in which a consumer sequentially considers unordered alternatives and the consumer search process is observed. Much of the empirical search literature focuses on dissimilar settings.⁴ First, De Los Santos et al. (2012) develop methods for estimating a fixed sample size search model whereas I develop methods for estimating a sequential search model. The sequential search model arguably better describes some settings, including contact lens e-commerce. Furthermore, although they analyze the same Comscore data that my article studies, they do not incorporate the panel dimension of these data in their analysis (they analyze books, which unlike contact lenses are typically purchased once). My article and De Los Santos et al. (2012) are complementary in that they provide empirical techniques for distinct sorts of search models that are differentially applicable to different settings. Another sort of non-sequential model in the literature is that of Allen et al. (2014), who develop a search model in which exerting search effort at a cost allows consumers to raise their chances of obtaining additional mortgage quotes. Turning to sequential search models, Koulayev (2014) models a consumer clicking through pages of hotel listings on a booking platform. This model is tailored to a

³Although Hortaçsu and Syverson (2004) allow for vertical (but not horizontal) differentiation between product offerings, Hong and Shum (2006) and Moraga-González and Wildenbeest (2008) use a model without seller differentiation.

⁴See Honka et al. (2019) for an overview of the empirical consumer search literature.

context in which alternatives are ordered, whereas my approach applies to contexts with an unordered set of retailers. Ursu (2018) similarly studies search of ordered hotel listings. Jolivet and Turon (2019) study sequential search for products within an e-commerce platform, although their approach is tailored to the case in which the consumer search process is not observed.

Three other literatures are relevant to my work. First, it relates to a literature that studies sources of limited consideration and market power in brick-and-mortar retail; see Sorensen (2000) for analysis of pharmacies and Dubois and Perrone (2015) for analysis of supermarkets. Second, it relates to a literature on inertia in consumer choice (including, e.g., Heckman 1981 and Kasahara and Shimotsu 2009), especially Dubé et al. (2009) and Dubé et al. (2010). Last, this article relates to a literature on platform design in ecommerce, including Dinerstein et al. (2018), who study search within eBay, and Lee and Musolff (2021) who study the interaction of seller differentiation and platform design on Amazon's Marketplace platform.

2 Setting and data

This study's primary data source is the Comscore Web Behavior Panel for 2007–2008 (Comscore 2007–2008). This dataset includes online browsing and transactions activities for a panel of US households.⁵ As noted by De Los Santos et al. (2012) and Saruya and Sullivan (2023), the Comscore panel is representative of online US consumers. The browsing data include a record of each web domain visited by a panelist; each record includes a panelist identifier and transactions associated with the visit.⁶ For each transaction, I observe the price and quantity of each purchased product.

The contact lens transactions analyzed in this article occur at three retailers that collectively account for about 95% of contact lens transactions in the

 $^{^{5}}$ The 2007 and 2008 panels include about 92 000 and 58 000 households, respectively.

⁶The data do not include the list of pages visited by a panelist within a web domain; for example, a record of a panelist visiting **amazon.com** does not reveal the visited product pages within Amazon.

data: 1-800 Contacts (1800), Vision Direct (VD), and Walmart (WM). The specialty retailers 1800 and VD in turn account for about 95% of sales among these three retailers. The former, 1800, launched in 1995 and was the market leader during 2007—2008 with a market share of about two-thirds. Vision Direct launched later, in 2004. Contact lens e-commerce was sizeable by 2007; 1800 made net sales of \$125 million in the first half of 2007. Although many new retailers have entered contact lens e-commerce since 2008, 1800 remained the market leader for many subsequent years.⁷

For each retailer and each brand of lenses, I construct a daily price time series. In doing so, I assume that the brand's price remains fixed at its most recent observed transaction price until the time of the subsequent observed transaction. This procedure introduces some measurement error, but the error is likely to be small because prices are updated often: the mean gap between transactions for top brands is generally under two weeks. The prices in the time series do not include shipping fees, although 1800 and VD both waived shipping fees for sufficiently large purchases.⁸

The dataset used in the article's analysis is a panel of search efforts, each of which is a sequence of store visits and a purchase decision. The purchasing alternatives here are visited stores and the outside option of not buying online. I construct the search effort for a transaction by determining all visits to retailers nearby in time to the transaction. Appendix O.2 details the procedure. For online retailers that exclusively sell contact lenses and associated products, there is little danger of incorrectly assuming that a consumer's visit to the retailer involved searching for contact lenses rather

⁷A response by 1800 to a Federal Trade Commission complaint (FTC Matter 141 0200, docket no. 9372, "Respondent 1-800 Contacts, Inc.'s Proposed Findings of Fact and Conclusions of Law") in 2017 stated that 1800 accounted for about 10% of total US contact lens sales whereas all purely online contact lens retailers accounted for about 17% of sales, implying that 1800's market share among online contact lens retailers was about 60%.

⁸1800, for example, offered free shipping on orders over \$50.

than some other product. ⁹ For other retailers and product categories, this is a risk than can be managed by obtaining data on within-website consumer search that identifies the product-specific webpages that the consumer visited within the website.

In the United States, optometrists and ophthalmologists prescribe contact lenses to their patients. A prescription specifies a brand, parameters (e.g., diameter and power), and an expiration date (typically one or two years in the future). I infer consumers' prescription based on the brand of lenses that they buy. When a consumer buys a different brand than that previously purchased, I assume the consumer's prescription has changed and that the consumer holds the new prescription alone until the next purchase. About 15% of consumers in the sample switch brands.¹⁰

Tables 1a and 1b describe consumer search efforts and retailers in the data. Table 1a describes consumer search efforts in the sample. The median price paid for a box of lenses was about \$30 and the median number of boxes of lenses purchased was two (one for each eye). Consumers make 2.5 search efforts on average, yet some consumers make many more search efforts. Table 1b reports the number of transactions at each retailer and each retailer's average relative price, defined as the across-transaction mean ratio of the retailer's price to 1800's price for the transacted brand at the time of the transaction. The table shows that 1800 had a market share of about two-thirds while charging higher prices than its rivals. VD had a market share of about 30% and offered contacts at 85% of 1800's prices, on average.

⁹Walmart, however, sold many other products online. Consequently, I use a more restrictive rule for including visits to Walmart in the sample (see Appendix O.2). I also demonstrate the robustness of model estimates to the treatment of Walmart (i.e., to dropping Walmart or treating it in the same way as the other retailers in constructing search efforts); see Online Appendix O.8.

¹⁰Online Appendix Tables O.6 and O.7 provide analysis of these switchers. One finding is that household size does not predict switching, which suggests that switching does not reflect distinct household members ordering different brands. In addition, the difference in mean prices faced by a consumer before and after switching is small and not statistically significant. One explanation is that switches are not driven by price considerations.

(a) Search efforts and transactions				(b) Retailers		
	Moon	Quantiles			N.	Mean
	Mean	0.25	0.75	Store	trans.	rel. price
Transaction price	31.05	19.99	38.95	1800	849	1.00
Transaction quantity	2.83	1.00	4.00	VD	416	0.85
N. search efforts	2.47	1.00	3.00	WM	70	0.94
N. transactions	1.65	1.00	2.00			
N. consumers $= 793$						
N. search efforts $= 1956$						
N. transactions $= 1310$						

Table 1: Descriptive statistics

Note: Table 1a reports descriptive statistics for the sample of search efforts and transactions, pooled across retailers. Table 1b reports descriptive statistics for each retailer. "N. trans." provides the numbers of transactions at each retailer in the sample. "Mean rel. price" reports the average ratio of the store's price to 1800's price across transactions.

3 **Descriptive analysis**

This section first provides evidence that consumers conduct little search online and often pay above the minimum available price in online markets for undifferentiated products. It then characterizes the influence of prices on consumer browsing and purchasing decisions.

Limited consideration 3.1

Active consideration of online retailers is severely limited in contact lens ecommerce. Table 2 displays the share of contact lens search efforts involving one, two, and three store visits. The "Baseline" column provides results for search efforts constructed from visits to 1800 or VD up to 14 days before a purchase as described in Appendix O.2. The "2 days before" column only includes visits made up to two days before a purchase or another visit. Under the baseline data construction, 83% of search efforts involve a visit to only one store. Table 2 also shows that search efforts are insensitive to the time-window used in constructing search efforts.

Consumers visit few stores despite the possibility of saving on lenses by

visiting and purchasing from other stores. Table 3 shows that 70% of transactions occur at a store that sells the purchased brand above the minimum price offered among the three major retailers. The magnitude of spending in excess of these minimum prices is significant — consumers pay, on average, 16.3% above the minimum available price. Additionally, in 43% of search efforts with multiple visits, the consumer does not choose the store with the lowest price among visited sites. On average, the consumer pays 7.1% over the minimum available price among visited sites. Search frictions provide one explanation for purchasing above the minimum available price. An alternative explanation is that some retailers offer superior customer service or shipping, and some consumers prefer to purchase from these retailers over lower-price rivals. The fact that 1800 outsells VD despite charging higher prices—see Table 1b—suggests that 1800 is more appealing to consumers in non-price dimensions. Alternatively, consumer awareness of 1800 could be higher than that of VD. The availability of both search and purchase data will allow me to distinguish between these explanations.

Table 2: Share of search efforts by number of visited store

# of	Share	of sessions
visits	Baseline	2 days before
1	0.83	0.84
2	0.16	0.15
3	0.01	0.01

Table 3: Transactions above minimum available price

	Value
Share of transactions above min price	0.70
Average payment over min price $(\$)$	4.31
Average payment over min price $(\%)$	16.3

Notes: this table reports the (i) share of transactions made above the minimum available price (MAP); (ii) the average difference of paid price and the MAP, and (iii) the average relative difference of the transaction price over the MAP.

3.2 Prices, browsing, and purchasing

I now turn to the role of prices in directing consumer behaviour. That 1800 boasted the highest sales despite charging the highest average prices could reflect that consumers generally prefer 1800, which could lead 1800 to charge higher prices. My solution to this price endogeneity problem in demand estimation is to exploit cross-brand differences in stores' relative prices. If stores' quality differences equally affect their sales of all brands, then the extent to which a store has relatively lower sales for brands that it sells for relatively higher prices is informative about consumer price sensitivity.¹¹.

To exploit between-brand variation to estimate price sensitivity, I specify store fixed effects in consumer utilities. I assess the suitability of this approach with descriptive multinomial logit regressions with and without fixed effects. An additional purpose of these regressions is to determine whether prices guide search, which would suggest that consumers have some knowledge of prices prior to search. The estimating equation is

$$u_{ift} = q_{ft} - \alpha p_{ift} + \varepsilon_{ift}, \qquad f \in \{1800, \text{WM}, \text{VD}\}, \tag{1}$$

where $y_{it} = \arg \max_{f} u_{ift}$ is either the store from which the consumer purchases or the first-visited store in a search effort, *i* indexes consumers, *t* indexes search efforts, and p_{ift} is retailer *f*'s price for *i*'s brand. Additionally, ε_{ift} is an unobservable iid type 1 extreme value (T1EV) shock. I estimate a specification without fixed effects in which $q_{ft} = \bar{q}$ for all *f* and *t* and one with fixed effects in which $q_{ft} = q_{f\tau}$, where τ indicates the half-year (e.g. first half of 2007). I estimate the regressions with the purchase decision as the outcome on a dataset of all search efforts that end in a transaction. I use a disjoint dataset of all search efforts that do not end in a transaction for the regressions with first-visited store as the outcome.

¹¹Online Appendix Figure O.1 shows how 1800 has relatively low sales among the brands for which it charges especially high prices relative to VD.

Specification 1: $q_{ft} = \bar{q} \forall f, t$		Specification 2: seller/half-year fixe effects			
	Purchase	First visit		Purchase	First visit
α	-0.006	-0.056		0.035	0.025
	(0.003)	(0.010)	α	(0.004)	(0.014)
olasticity	-0.072	-0.092	Average	0.449	0.455
elasticity	(0.043)	(0.080)	elasticity	(0.049)	(0.111)

Table 4: Descriptive multinomial logit regressions (contact lenses)

Notes: The table reports maximum likelihood estimates of (1) for the contact lenses category. Standard errors are reported in parentheses. The "Average elasticity" is the average own-price elasticity taken across transactions.

Table 4 reports results. Without fixed effects, I estimate that consumers are more likely to purchase from sellers charging higher prices. This relationship is reversed upon the introduction of fixed effects. Additionally, the first-visited store responds to prices in a similar way as purchases. This suggests that consumers have some knowledge of prices before search.¹²

4 Model of consumer search

This section outlines a general model of repeated sequential search. Each consumer *i* searches for a product *j* across *F* retailers at different occasions in time. The consumer makes search efforts $t \in \{1, \ldots, T_i\}$ at exogenously determined times. In each effort, the consumer determines which retailers $f \in \mathcal{F} = \{1, \ldots, F\}$ to visit. Retailer *f* charges a price p_{ift} for consumer *i*'s brand during a search effort *t*. The consumer additionally chooses a store *f* among visited stores from which to purchase, or not to buy product *j* online (denote f = 0). The consumer incurs a search cost κ_{ift} for visiting retailer *f* in search effort *t*. In the context of e-commerce, search costs capture costs of learning about store *f*, navigating to the webpage on which it lists product

¹²The elasticities in Specification 2 of Table 4 fall below one, which is inconsistent with profit maximization. The low values of the estimated average elasticities may reflect misspecification of the simple logit demand model, which does not capture horizontal differentiation of retailers, search frictions, or state dependence. The estimates of the preferred model imply more reasonable elasticities; see Online Appendix Table O.4.

j for sale, and determining whether product j is in stock; these costs may be influenced by display, search-engine, or other advertisements. Consumers conduct sequential search according to the optimal strategy of Weitzman (1979). Consumer *i*'s utility from purchasing from store f during search effort t is

(Online)
$$u_{ijft} = q_f - \alpha_i p_{ift} + \phi h_{ift} + \gamma_{if} + \varepsilon_{ift}$$
 (2)

(Offline)
$$u_{ij0t} = \varepsilon_{i0t},$$
 (3)

where q_f governs the quality of store f; γ_{if} is consumer *i*'s persistent taste for f; ε_{ift} is consumer *i*'s *t*-specific match value with f. Store quality here captures shipping speed, customer service, returns policies, and persuasive effects of advertising. Additionally, h_{ift} is an indicator for whether the consumer purchased from f in search effort t - 1. I refer to $\mathbf{h}_{it} = \{h_{ift}\}_{f \in \mathcal{F}}$ as consumer *i*'s *state*. Price sensitivity depends on $\alpha_i = \alpha_0 + \alpha_1 I_i$, where I_i is an indicator for consumer *i*'s household income exceeding \$75,000. Additionally, ϕ governs state dependence, which may arise from habit formation, switching costs, or store loyalty.

Before search, the consumer knows all but the ε_{ift} match values. Section 4.3 justifies this assumption. I also assume that consumers are myopic in not anticipating the effects of their choices on future payoffs, a common assumption in the state dependence literature (e.g., Dubé et al. 2010).

The optimal sequential search strategy of Weitzman (1979) involves visiting stores in descending order by *reservation utility* until obtaining an indirect utility higher than the maximum reservation utility among unsearched stores. Consumer *i*'s reservation utility r_{ift} for store *f* in search effort *t* is

$$\kappa_{ift} = \int_{r_{ift}}^{\infty} (u - r_{ift}) dF_{ift}(u), \qquad (4)$$

where κ_{ift} is the search cost that consumer *i* incurs for visiting store *f* in search effort *t* and $u_{ift} \sim F_{ift}$ conditional on all but ε_{ift} . Note that, by

construction, the consumer is indifferent between (i) enjoying a payoff of r_{ift} without further search and (ii) visiting store f before enjoying a payoff equal to the maximum of u_{ift} and r_{ift} . Reservation utilities can be written as

$$r_{ift} = q_f - \alpha_i p_{ift} + \phi h_{ift} + \gamma_{if} + \Gamma_0^{-1}(\kappa_{ift}), \qquad (5)$$

for $\Gamma_0(\kappa) = \int_{\kappa}^{\infty} (u-\kappa) dF_0(u)$, where F_0 is the distribution of the ε_{ift} match values, assumed T1EV. Because Γ_0 and its inverse are strictly decreasing functions, a store's reservation utility is decreasing in its search cost. Kim et al. (2010) and Moraga-González et al. (2023) similarly invert equations defining reservation utilities to obtain expressions resembling (5).

There is a convenient parametric distribution of the search costs κ_{ift} that yields tractable choice probabilities for search effort outcomes. Suppose that $\kappa_{ift} \sim F_{\kappa}(\cdot; \bar{\kappa}_f)$ independently of all else (including search costs for other search efforts $t' \neq t$ or other consumers i), where

$$F_{\kappa}(\kappa;\bar{\kappa}_f) = 1 - \exp\left\{-\exp\left\{-\Gamma_0^{-1}(\kappa) - \bar{\kappa}_f\right\}\right\}.$$
(6)

The $\bar{\kappa}_f$ parameter positively relates with both the mean and variance of the distribution of search costs for store f. Differences in this parameter across retailers f reflects differences in awareness of retailers and in ease of navigating to and within retailers' websites. Figure 1 plots $F_{\kappa}(\cdot; \bar{\kappa})$. Under this distribution, we can express equation (5) as

$$r_{ift} = q_f + \gamma_{if} - \alpha_i p_{ft} + \phi h_{ift} - \bar{\kappa}_f + \eta_{ift}, \tag{7}$$

where the η_{ift} are mutually independent (across *i*, *f*, and *t*) T1EV random variables. To see why, note that

$$\Pr(\Gamma_0^{-1}(\kappa_{ift}) \le x) = \Pr(\kappa_{ift} \ge \Gamma_0(x))$$

$$= 1 - F_{\kappa}(\Gamma_0(x), \bar{\kappa}_f)$$

$$= \exp\left\{-\exp\left\{-\Gamma_0^{-1}(\Gamma_0(x)) - \bar{\kappa}_f\right\}\right\}$$

$$= \exp\left\{-\exp\left\{-\exp\left\{-(x + \bar{\kappa}_f)\right\}\right\},$$
(8)

which is the distribution function of a T1EV random variable with location





parameter $-\bar{\kappa}_f$. Thus, $\Gamma_0^{-1}(\kappa_{ift}) + \bar{\kappa}_f \sim \eta_{ift}$, where η_{ift} is a standard T1EV random variable. Substituting $\Gamma_0^{-1}(\kappa_{ift})$ for $-\bar{\kappa}_f + \eta_{ift}$ in (5) yields (7). Note that, by assuming that search costs are independent across search efforts, I rule out a dependence of search costs in search effort t on consumer search behaviour in previous search efforts t' < t.

The distribution above is one of the two model features that give rise to tractable choice probabilities. The other is a bijective mapping between (i) inequalities relating reservation and indirect utilities and (ii) outcomes of search efforts. Given the distributional assumptions, these inequalities yield closed-form outcome probabilities. To illustrate, suppose that a consumer visits stores f and f' before buying from f. This sequence of visits implies that the highest reservation utility is that for f and that the reservation utility for f' exceeds the indirect utility for store f. Otherwise, the consumer would have terminated search after visiting f to buy from that store. Analogous reasoning establishes that the reservation utility for f' exceeds u_{i0} . Because the consumer purchases from f, the indirect utility of f must exceed the indirect utilities. This reasoning is summarized by the following chain of inequalities (wherein I suppress the t subscript):¹³

$$r_{if} \ge r_{if'} \ge u_{if} \ge u_{i0} \lor u_{if'} \lor \max_{g \in \mathcal{F} \setminus \{f, f'\}} r_{ig}$$

¹³Note that \lor is the maximum operator, i.e. $a \lor b = \max\{a, b\}$.

The probability that r_{if} exceeds $r_{if'}, u_{if}, u_{i0}, u_{if'}$, and $\max_{g \in \mathcal{F} \setminus \{f, f'\}} r_{ig}$ takes the standard logit form:

$$\frac{e^{\bar{r}_{if}}}{\sum_{g=1}^{F} e^{\bar{r}_{ig}} + e^{\bar{u}_{i0}} + e^{\bar{u}_{if}} + e^{\bar{u}_{if'}}},$$

where $\bar{u}_{ig} = u_{ig} - \varepsilon_{ig}$ and $\bar{r}_{ig} = r_{ig} - \eta_{ig}$. Similarly, the probability that $r_{if'}$ exceeds u_{if} , u_{i0} , $u_{if'}$, and $\max_{g \in \mathcal{F} \setminus \{f, f'\}} r_{ig}$ also has a standard logit form,

$$\frac{e^{\bar{r}_{if'}}}{\sum_{g \neq f} e^{\bar{r}_{ig}} + e^{\bar{u}_{i0}} + e^{\bar{u}_{if}} + e^{\bar{u}_{if'}}}$$

as does the probability that u_{if} exceeds u_{i0} , $u_{if'}$, and $\max_{g \in \mathcal{F} \setminus \{f, f'\}} r_{ig}$:

$$\frac{e^{\bar{u}_{if}}}{\sum_{g\notin\{f,f'\}}e^{\bar{r}_{ig}}+e^{\bar{u}_{i0}}+e^{\bar{u}_{if}}+e^{\bar{u}_{if'}}}$$

Given the independence of irrelevant alternatives property of the logit, we then obtain the overall probability of the search outcome by multiplying together the probabilities above:

$$\frac{e^{\bar{r}_{if}}}{\sum_{g=1}^{F} e^{\bar{r}_{ig}} + e^{\bar{u}_{i0}} + e^{\bar{u}_{if}} + e^{\bar{u}_{if'}}} \times \frac{e^{\bar{r}_{if'}}}{\sum_{g \neq f} e^{\bar{r}_{ig}} + e^{\bar{u}_{i0}} + e^{\bar{u}_{if}} + e^{\bar{u}_{if'}}}}{e^{\bar{u}_{if}}} \times \frac{e^{\bar{u}_{if}}}{\sum_{g \notin \{f,f'\}} e^{\bar{r}_{ig}} + e^{\bar{u}_{i0}} + e^{\bar{u}_{if}} + e^{\bar{u}_{if'}}}}.$$
(9)

Online Appendix O.1 provides the inequalities corresponding to other outcomes under an arbitrary number of alternatives.

4.1 Discussion of results

The choice probabilities in (9) are straightforward to compute. Without using either the search cost distribution (6) or the chains of inequalities implied by the Weitzman (1979) strategy, computing choice probabilities would require, for a given draw of unobservables κ_{ift} and ε_{ift} , the inversion of a function defined by an integral (i.e., Γ_0) to compute reservation utilities. It would then require the sequential solution of the consumer's search problem by comparing reservation and indirect utilities at each step in search. Last, it would require integration over κ_{ift} and ε_{ift} to obtain choice probabilities. The mapping between chains of inequalities and search effort outcomes reduces the burden of computing choice probabilities even under arbitrary dependence structures of $(\varepsilon_{ift}, \kappa_{ift})$ across consumers, time, and retailers.¹⁴ Indeed, one could simulate search efforts by drawing a sequence $\{(\varepsilon_{ift}, \kappa_{ift})\}_{f \in F}$ of random variables from an arbitrary distribution and then determine the associated outcome of search by identifying the set of inequalities satisfied by this sequence. This procedure facilitates estimation of the model using simulation-based estimators such as the indirect inference estimator used in this article (see Section 6.1). Notably, this procedure allows the assumption that search costs are iid across time to be relaxed.

Although the article's parametric restrictions are not necessary for tractable analysis of the model, they simplify computation in several ways. First, they facilitate maximum likelihood estimation.¹⁵ This is because the parametric restrictions yield exact closed-form choice probabilities. Without these exact closed forms, the researcher must approximate choice probabilities using computational methods such as simulation in order to compute model likelihoods. In addition, the parametric restrictions simplify the simulation of search efforts. When the researcher uses these parametric restrictions, it is possible to simulate search efforts by either (i) assessing which inequalities characterizing search effort outcomes hold under a given unobservable draw $\{(\varepsilon_{ift}, \kappa_{ift})\}_{f \in F}$ or (ii) drawing directly from simple closed-form choice probabilities. Option (i) involves assessing many pairwise inequalities, which makes option (ii) more convenient in general. Last, the closed-form expressions are differentiable with respect to model parameters, which facilitates

¹⁴Other articles have exploited utility rankings in analyzing search models. Moraga-González et al. (2023) specify inequalities based on a result of Armstrong (2017) and Choi et al. (2018). Morozov et al. (2021) and Ursu (2018) pool separate inequalities for (i) visit order, (ii) stopping decision, and (iii) purchase decision. Ursu et al. (2023) describe various methods for simulating search effort outcomes using distinct inequalities for visit order, stopping, and purchase decision derived from the Weitzman search strategy.

¹⁵As noted in Section 6, I find that indirect-inference estimators are better behaved in my setting than maximum likelihood estimators. With that said, maximum-likelihood estimators boast greater asymptotically efficiency than indirect-inference estimators and thus may be more appropriate in other sequential search settings.

use of derivative-based optimization procedures in estimation and the computation of demand elasticities in retailer pricing analysis.

Note that this article's framework is readily applied to the analysis of search across products within a retailer or e-commerce platform; this simply requires re-labelling sellers f as products. It is also possible to use the article's mapping between utility inequalities and search effort outcomes and the proposed search cost distribution to analyze a model of search over both retailers and products offered by each retailer (e.g., over both book titles and bookstores). Estimating such a model, though, would require data on both across- and within-retailer search. Upon specifying a within-retailer search problem, the researcher could enter the inclusive value of the within-retailer problem as a term in the store-level indirect utilities that the consumer considers in across-retailer search. When the search model described here is applied to both search problems, analyses of choice probabilities at each stage of consumer search would be facilitated by this article's methods.

4.2 Probabilities of sequences of search efforts

Search efforts at different times are related by state dependence and persistent tastes. In this section, I provide an expression for the probability of a consumer's sequence of search efforts across time. Let $\boldsymbol{y}_i = \{y_{it}\}_{t=1}^{T_i}$, where y_{it} denotes consumer *i*'s search/purchase choices in search effort *t*. Similarly let $\boldsymbol{p}_i = \{\boldsymbol{p}_{it}\}_{t=1}^{T_i}$, where $\boldsymbol{p}_{it} = \{p_{ift}\}_{f\in\mathcal{F}}$ denotes the prices of consumer *i*'s brand at search effort *t* across retailers *f*. Next, let \boldsymbol{h}_{i1} denote consumer *i*'s initial state, let $\boldsymbol{\gamma}_i = \{\gamma_{if}\}_{f\in\mathcal{F}}$ denote the consumer's persistent unobserved tastes, let $\boldsymbol{\theta}$ denote an arbitrary parameter vector, and let $\boldsymbol{\theta}_0$ denote the true parameter vector. The model provides conditional probabilities of search effort outcomes that I denote by $\Pr(y_{it}|I_i, \boldsymbol{p}_{it}, \boldsymbol{h}_{it}, \boldsymbol{\gamma}_i; \boldsymbol{\theta})$. The overall

conditional probability of consumer i's sequence of search efforts

$$\Pr(\boldsymbol{y}_i \mid I_i, \boldsymbol{p}_i, \boldsymbol{h}_{i1}; \theta) = \int \Pr(\boldsymbol{y}_i \mid I_i, \boldsymbol{p}_i, \boldsymbol{h}_{i1}, \boldsymbol{\gamma}_i; \theta) dG(\boldsymbol{\gamma}_i \mid \boldsymbol{p}_i, \boldsymbol{h}_{i1}; \theta),$$

where G is the distribution of $\boldsymbol{\gamma}_i$ conditional on \boldsymbol{p}_i and \boldsymbol{h}_{i1} .

Two econometric problems arise when integrating over γ_i . The first is the standard initial condition problem: the distribution of $\boldsymbol{\gamma}_i$ conditional on \boldsymbol{p}_i and \boldsymbol{h}_{i1} will depend on \boldsymbol{h}_{i1} because \boldsymbol{h}_{i1} reflects consumers' past choices, which depended on γ_i . Thus, we cannot drop h_{i1} from the conditioning set. The second problem, which I call the endogeneity problem, relates to the dependence of $\boldsymbol{\gamma}_i$ and prices \boldsymbol{p}_i conditional on \boldsymbol{h}_{i1} . To understand this dependence, suppose that store f sold two products and that its price for the first product was high relative to other stores whereas its price for the second product was relatively low. In that case, consumers seeking the first product who buy at f require favourable tastes for the store to outweigh f's high price. Similarly, consumers seeking the second product may buy from f despite disliking the store due to its low price. Thus, the prices faced by a consumer and the consumer's tastes for stores are generally correlated conditional on the initial state. Online Appendix O.5 presents evidence that consumers who previously purchased contact lenses from a high-price seller especially like that seller.

The problems noted above invalidate the simplifying assumption that $G(\gamma_i | p_i, h_{i1}; \theta)$ depends neither on the initial state nor on prices. I address these problems by specifying a parametric model of γ_i 's conditional distribution:

$$\gamma_{if} \mid (\boldsymbol{p}_i, \boldsymbol{h}_{i1}) \sim \begin{cases} N\left(\lambda \tilde{p}_{if}, \sigma_{\gamma}^2\right), & h_{if1} = 1\\ N\left(\Gamma_{fg}, \sigma_{\gamma}^2\right), & h_{ig1} = 1 \end{cases}$$
(10)

where g denotes a seller other than f; λ , Γ_{fg} , and σ_{γ}^2 are parameters; and \tilde{p}_{if}

is the relative price of consumer i's brand at f at i's first observed purchase:

$$\tilde{p}_{if} = \left(p_{if1} - \frac{1}{F}\sum_{g=1}^{F} p_{ig1}\right) / \frac{1}{F}\sum_{g=1}^{F} p_{ig1}.$$

The parameter λ governs the extent to which consumers who initially buy from f despite its high price have more favourable tastes for f. The parameter Γ_{fg} governs the tastes for store f of consumers who initially buy from store g. Last, σ_{γ}^2 governs variability in persistent store tastes.

My approach to modelling γ_i is based on commonly used approaches in panel data settings. Specifying a parametric distribution of γ_i conditional on the initial state follows Wooldridge (2005). Also, modelling the dependence of γ_i on prices conditional on the initial state follows the correlated random effects (CRE) approach (Chamberlain 1980, Mundlak 1978, Wooldridge 2010), which involves modelling the dependence of unobservables on regressors.

4.3 Justification of assuming search over match value

The assumption of known prices and search over match values is common in the consumer search literature (e.g., Kim et al. 2010, Moraga-González et al. 2023). It is justified in my context for several reasons. First, regressions from Section 3.2 suggest that consumers respond to prices in choosing stores to visit even when they do not ultimately buy lenses. This is compatible with the consumer choosing visits based on knowledge of prices. Consumers may know prices based on previous search experience—recall that I drop consumers' first search efforts from the sample—or through adverts.¹⁶ Another reason to assume search over match values is the presence of non-price retailer characteristics that consumers learn through search. These include the consumer's perception of the retailer's website usability and design; timevarying marketing materials on retailers' websites; the speed at which the

¹⁶This is plausible given that 1800 advertised heavily in the sample period, with advertising expenses equal to 12% of costs of goods sold in the first half of 2007.

retailer can verify prescriptions (which may depend on the retailer's website traffic and on the consumer's internet speed); and time-varying public reviews of retailers that consumers may differentially discover in the search process. Perhaps the most important source of variation in match values relates to prescription/brand-specific inventories and shipping times. Contact lenses vary not only by brand but also by other prescription parameters; these include base curve, power, sphere, etc.¹⁷ Whether a retailer has a specification in stock determines the store's shipping time for an order. This likely explains why 1800's advertisements boasted of the firm's large inventories. Furthermore, a response by 1800 in 2017 to an FTC complaint suggests the importance of inventory in contact lens retail: it claims that a consumer could wait 4–8 weeks for a shipment from a rival online retailer if the retailer did not have the consumer's prescription in stock and that independent eye-care professionals typically had only about 40% of orders in stock.¹⁸

An alternative approach is to assume search over prices and specify consumer beliefs over prices.¹⁹ If consumer beliefs concentrate around the true prices on account of the common rational expectations assumption, this approach is similar to one that assumes knowledge of prices but fails to account for non-price information uncovered by search. A general difference between models of search over match values and search over prices is that, in the latter, search costs attenuate consumer responses to price. This is because consumers do not condition their choice of store to visit on price (given that they do not know stores' prices prior to search). Furthermore, they may not visit other stores even upon finding a high price at a visited store on account of search costs, thus accepting higher prices than they would

¹⁷Prices do not vary by these parameters.

¹⁸See FTC Matter 141 0200, docket no. 9372, "Respondent 1-800 Contacts, Inc.'s Proposed Findings of Fact and Conclusions of Law."

¹⁹See Mehta et al. (2003), Hong and Shum (2006), Moraga-González and Wildenbeest (2008), and Honka (2014).

absent search costs. In models of search over match values, search costs do not necessarily blunt consumer price sensitivity and may in fact amplify it. Indeed, Choi et al. (2018)—who argue for the relevance of models of search over match values for e-commerce—show that search costs may lower prices in oligopolistic competition in such models. This reflects that the sensitivity of search to prices may amplify the overall response of sales to prices. The divergence between the search models above is relevant, e.g., for analysis of the effect of search costs on the mean payment over the minimum available price. In reality, consumers likely search over both price and nonprice characteristics, and I choose to model search over the latter based on features of the setting under investigation as enumerated above. I leave study of the article's sensitivity to the choice of search model to future research.

5 Price competition

To analyze market power, I specify a pricing model. Although the model is static in that each retailer sets a time-invariant price for each product, the model captures long-run responses of consumer states to prices. An alternative approach is to study Markov perfect equilibria (MPE) of a dynamic pricing game. Whereas it is straightforward to find Nash equilibria of the static model, solving for MPE requires model simplifications given the infinite dimensionality of the state space. I analyze a dynamic model that is simplified in two main ways: Walmart is excluded from the model and the distribution of γ_i is approximated by a discrete distribution with two support points. See Online Appendix O.6 for details. The dynamic model

A challenge in modelling static pricing is accounting for state dependence in demand. I propose a *long-run demand* system that represents consumer choice under the long-run distribution of states. This system involves *long-* run state probabilities $\{\rho_f(\boldsymbol{p},\boldsymbol{\gamma}_i,\alpha_i)\}_{f=1}^F$, defined as the solutions of

$$\rho_f(\boldsymbol{p}, \boldsymbol{\gamma}_i, \alpha_i) = \sum_g \sigma_{fg}(\boldsymbol{p}, \boldsymbol{\gamma}_i, \alpha_i) \rho_g(\boldsymbol{p}, \boldsymbol{\gamma}_i, \alpha_i) \qquad \forall f, \tag{11}$$

where $\sigma_{fg}(\mathbf{p}, \mathbf{\gamma}_i)$ is the probability with which a consumer with state $h_{igt} = 1$ buys from store f given prices p. The right-hand side of (11) is the overall probability of a consumer belonging to state f after a search effort when the probability that consumer belongs to state g prior to search is $\rho_g(\mathbf{p}, \mathbf{\gamma}_i, \alpha_i)$. Thus, condition (11) imposes that the share of type- $(\mathbf{\gamma}_i, \alpha_i)$ consumers in state f is stable. Letting H denote the unconditional distribution of $(\mathbf{\gamma}_i, \alpha_i)$, the long-run market share for store f is

$$\sigma_f^L(\boldsymbol{p}) \coloneqq \int \sum_g \rho_g(\boldsymbol{p}, \boldsymbol{\gamma}_i, \alpha_i) \sigma_{fg}(\boldsymbol{p}, \boldsymbol{\gamma}_i, \alpha_i) dH(\boldsymbol{\gamma}_i, \alpha_i).$$

6 Estimation

6.1 Indirect inference

I estimate the model using an indirect inference (I-I) estimator.²⁰ This approach involves (i) computing auxiliary statistics $\hat{\beta}_n$ on the sample; (ii) simulating outcomes under a trial model parameter value θ ; and (iii) computing the statistics on the simulated data, letting $\tilde{\beta}_n(\theta)$ denote the statistics computed on the simulated data. The I-I estimator $\hat{\theta}$ minimizes a measure of the distance between $\hat{\beta}_n$ and $\tilde{\beta}_n(\hat{\theta})$:

$$\hat{\theta}_n = \arg\min_{\theta} \left(\hat{\beta}_n - \tilde{\beta}_n^H(\theta) \right)' \hat{\Omega}_n \left(\hat{\beta}_n - \tilde{\beta}_n^H(\theta) \right)$$
(12)

where $\hat{\beta}_n$ are ordinary least squares (OLS) estimators computed on the sample and $\tilde{\beta}_n^H(\theta)$ are the same OLS estimators computed on outcomes simulated under θ conditional on $\{\boldsymbol{x}_i, \boldsymbol{h}_{i1}\}_i$, where \boldsymbol{x}_i includes the consumer's brand and prices faced by the consumer. These outcomes are simulated H = 50times for each panelist. I simulate outcomes for a panelist by first drawing

²⁰See Gouriéroux et al. (1993). I use an I-I estimator instead of a maximum likelihood estimator (MLE) because that MLEs tend to exhibit poor finite-sample performance in discrete-choice settings with many low probability potential outcomes; see Krasnokutskaya and Seim (2011), Pakes et al. (2007), and Collard-Wexler (2013).

 γ_i under θ conditional on $\{x_i, h_{i1}\}$ according to the conditional joint distribution (10). To facilitate the treatment of past purchases as observable variables in studying state dependence, I drop each consumer's search efforts made before and including the consumer's initial purchase. To simulate the outcome of the first search effort following this initial search effort, I compute the probability of each search effort outcome using the closed-form expressions obtaining under the article's maintained parametric distributions. These probabilities allow for straightforward simulation of search effort outcomes. The simulated outcome of the search effort implies a state h_{i2} for the consumer's next search effort. I similarly simulate following search efforts. The $\hat{\Omega}_n$ object in (12) is a weighting matrix; I use the approximately optimal weighting matrix proposed in Online Appendix O.4. This appendix also provides the expression for the estimator's asymptotic variance on which I base inference.

The following list summarizes the regression coefficients included in $\hat{\beta}_n$.

- (i) Stores' visit shares: shares of search efforts with a visit to each store f.
- (ii) Stores' purchase shares: shares of search efforts with a purchase from each store f.
- (iii) Consideration set size: share of search efforts wherein the consumer visited all stores.
- (iv) Inertia share: share of search efforts with the same first-visited store as the associated consumer's previous search effort.
- (v) Inertia regression: regressions of indicators for whether a consumer visited a store on lagged purchases.
- (vi) Role of lagged price: regressions of an indicator for buying from 1800 on the contemporaneous and lagged price at 1800 (to target ϕ and γ_i related parameters).

- (vii) *Price sensitivity*: regression of purchase decisions on prices.
- (viii) Cross-visiting: the share of consumers who visit retailer f among those who last purchased from f', for each f, f' pair with $f \neq f'$.
 - (ix) Dependence of tastes and prices conditional on initial state: regressions of indicators for whether the consumer visited a particular store on the ratio of the store's price to the average price across stores.
 - (x) Price sensitivity heterogeneity: regression of transaction price relative to the minimum available price for the consumer's brand on an indicator for the consumer's household income exceeding \$75,000.

Appendix A details these statistics. It also reports their values on both the estimation sample and on data simulated from the model. Further, Online Appendix O.9 characterizes the sensitivity of the parameter estimates to the values $\hat{\beta}_n$ of the auxiliary statistics. The results in Online Appendix O.9 are consistent with the identification discussion in the proceeding subsection.

In estimation, I de-mean the prices that enter consumer utilities by the average price across stores conditional on brand and time. Without demeaning prices, the model would mechanically predict a larger probability of choosing the outside option for expensive brands.

6.2 Identification

The model features three groups of parameters: those affecting search costs $(\bar{\kappa}_f)$, those affecting consumers' purchasing utilities in a static fashion $(q_f, \alpha_0, \text{ and } \alpha_1)$, and those affecting consumers' purchasing utilities in a dynamic fashion (ϕ and the parameters governing the distribution of γ_i).²¹ Here, I

²¹Note that the assumption that search costs are iid across time implies that search costs do not affect consumer behaviour in a dynamic manner. With that said, it appears possible to separately identify persistence in search costs from persistence in unobserved tastes for stores based on their different implications for search versus purchase behaviour. The argument in the proceeding paragraph suggests why this is the case.

describe how parameters in each group are identified.

First consider the separate identification of parameters affecting search costs and those affecting purchasing utilities. The challenge here is that a store f's low sales could owe to either high costs of visiting store f or low consumer preferences for purchasing from store f (i.e., low indirect utilities u_{ift}). These two explanations are separately identified with data on the search process. Indeed, the extent to which each explanation holds is identified by the rate at which consumers who visit the store ultimately buy from the store. A store f having many visitors but few buyers indicates that it has low search costs but also low indirect utilities. Conversely, a store f having few visitors but a high rate of converting visitors into buyers indicates that it has high search costs but high indirect utilities.

Price endogeneity poses a challenge in identifying the static preference parameters. Here, price endogeneity arises from the fact that unobserved retailer quality influences retailer pricing. The first assumption that permits identification of the price coefficient parameters is that retailer quality does not vary across brands. This assumption permits the specification of retailer fixed effects that capture brand-invariant retailer quality. With these retailer fixed effects specified, the information that identifies the price coefficient is the covariance across brands between (i) stores' relative prices for a brand and (ii) stores' relative market shares for a brand.²² The assumption underlying the identification argument above would fail if retailers' quality varied across brands of contact lenses in a manner that correlated with prices. Given that return policy and customer service assurances on retailers' websites did not condition on the brand purchased, this sort of violation seems unlikely.

In my setting, identification of the price coefficient parameters also relies on

 $^{^{22}{\}rm Online}$ Appendix Figure O.1 describes this covariance.

the assumption that consumers cannot substitute across brands of lenses. The model explains the negative covariance between relative prices and relative market shares using substitution across stores by consumers with a fixed brand.²³ Another explanation for this covariance is substitution across brands. To illustrate, a consumer who enjoys 1800 may encourage their doctor to prescribe a brand that 1800 sells for a relatively low price. Such behaviour would also contribute to a negative covariance between a retailer's relative price for a brand and its relative market share in sales of that brand. By attributing the entirety of the covariance to within-brand substitution, I risk overstating price sensitivity α_i . With that said, the assumption of within-brand substitution is defensible: consumers do not have complete control over their brands as medical professionals ultimately prescribe brands, in part due to patients' optical needs.

Last, I discuss the identification of parameters affecting choice dynamics. The primary challenge here is the separate identification of state dependence and unobserved heterogeneity γ_{if} . Although both elements of preferences promote inertia, they have different empirical implications. Conditioning on a consumer, a model with switching costs features dependence of a consumer's choice on the previous choice whereas a model without switching costs does not. Additionally, in the context of my model, stronger persistent store tastes generate greater correlation between contemporaneous choice and choice two or more purchasing occasions ago conditional on the choice in the previous purchasing occasion than does strong state dependence. This motivates my inclusion of a regression of the consumer's contemporaneous choice on lagged choices among the I-I auxiliary statistics. One weakness of this approach is that I do not observe repeat purchasing by all consumers—

²³Given that a consumer is limited to choosing a single prescribed brand in the model, and all choice is between sellers of this one brand, brand fixed effects would not be identified; they would shift the attractiveness of all alternatives equally. For the same reason, unobserved brand characteristics—a usual source of price endogeneity—do not cause an identification problem in the model.

see Table 1a—and thus estimates obtained using the approach reflect the preferences of the subset of consumers who do search repeatedly. By assuming a constant state dependence parameter, I extrapolate the extent of state dependence found among repeated consumers to the entire population. The validity of this approach requires that state dependence does not systematically vary across groups of consumers who make different numbers of purchases in the data.

7 Parameter estimates

Table 5 reports parameter estimates. The "Baseline" panel reports results for the baseline model whereas the "Stripped down" panel reports results for a specification without state dependence or persistent heterogeneity. Under the baseline estimates, retailer-specific median search costs among consumers with household incomes under \$75,000 range from \$0.41 to \$1.29.²⁴ These median search costs are low, relative to the median transaction price of about \$30. In addition, search costs are lowest for 1800 and highest for VD, suggesting that 1800's sales advantage could owe to greater consumer awareness of 1800 relative to its rivals. The estimates suggest, however, that taste heterogeneity and state dependence exercise significant influence on consumer decisions: the σ_{γ}^2 parameter estimate indicates substantial dispersion in persistent tastes for retailers, the estimate of ϕ implies that having previously purchased from a store raises the consumer valuation of the store by \$4.48 for the median consumer, and the negative estimate of α_1 indicates that higher-income consumers are less price sensitive.

A comparison of the "Baseline" and "Stripped down" results suggests that ruling out state dependence and persistent heterogeneity leads to an overstatement of search costs. When these aspects of consumer are ignored, the

²⁴The $\bar{\kappa}_f$ parameters are not directly interpretable as search costs; they are parameters that govern the distribution of search costs. See Figure 1. This explains why the $\bar{\kappa}_f$ parameters may be negative without implying that search costs are negative.

model requires higher search costs to rationalize highly limited search.²⁵

Table 5b reports estimates of the mean consumer taste for retailer $f q_f + \mathbb{E}[\gamma_{if}]$, which I interpret as retailer quality. In line with 1800 selling more than its rivals at higher prices, 1800's estimated quality exceeds those of WM and VD. There are various reasons to expect that 1800 boasted higher quality than VD. In October 2007, 1800's website mentioned that 1800 employed 300 call centre representatives trained in ocular health and answered 90% of calls within 10 seconds. The website also stated that 1800 shipped 90% of orders within 24 hours, offered a "100% satisfaction guarantee" return policy, and accepted returns of unused lenses upon prescriptions changes. By contrast, VD's website in September 2007 did not describe customer service, shipping, or a return policy. A 2017 response by 1800 to a Federal Trade Commission complaint also suggested that offering a high quality of service was central to 1800's business strategy, whereas VD focused on offering lower prices.²⁶

The idiosyncratic tastes for retailers γ_{if} could reflect heterogeneity in tastes for the services that retailers differentially offer (e.g., quick shipping, generous return policies) or retailer marketing strategies targeted at specific consumer segments. Taste heterogeneity of this sort likely correlates with consumer characteristics. I find that consumer characteristics substantially explain purchase behaviour: a multinomial logistic regression of store of purchase on consumer characteristics yields a McFadden's R^2 of 0.23. Furthermore, the estimates suggest that consumers who have higher incomes, who have broadband, and who live in smaller households are more likely to purchase from 1800. Such consumer characteristics are determinants of the γ_{if} unobservables. Online Appendix Table O.3 and Figure O.2 detail the

²⁵The large estimates of median search costs, as well as the large standard errors for these estimates, reflect both higher estimated search costs (compare the $\bar{\kappa}_f$ parameter estimates) and lower estimates of price sensitivity α_0 .

 $^{^{26}\}mathrm{See}$ FTC Matter 141 0200, docket no. 9372, "Respondent 1-800 Contacts, Inc.'s Proposed Findings of Fact and Conclusions of Law."

regressions outlined above.

	Baseli	Baseline		Stripped down		Mean taste for f
Parameter	Estimate	SE	Estimate	SE	f	$Q_f = q_f + \mathbb{E}\gamma_{if}$
q_{1800}	-0.34	(0.11)	-0.10	(0.16)	1800	-1.22
$q_{ m WM}$	-2.23	(0.17)	-1.21	(0.25)	WM	-3.37
$q_{ m VD}$	0.30	(0.09)	0.24	(0.04)	VD	-3.67
ϕ	0.49	(0.13)	-	-		
$lpha_0$	0.11	(0.01)	0.02	(0.02)		
α_1	-0.08	(0.04)	-	-		
$\bar{\kappa}_{1800}$	-2.71	(0.35)	-0.31	(0.30)		
$\bar{\kappa}_{ m WM}$	-1.89	(0.14)	0.40	(0.23)		
$\bar{\kappa}_{ m VD}$	-1.55	(0.22)	0.92	(0.13)		
$\Gamma_{1800,VD}$	-3.26	(0.43)	-	-		
$\Gamma_{\rm VD,1800}$	-5.57	(1.08)	-	-		
σ_{γ}^2	1.30	(0.17)	-	-		
$\lambda^{'}$	3.99	(1.57)	-	-		
Med. SC						
1800	0.41	(0.17)	28.19	(47.34)		
WM	0.93	(0.14)	51.32	(77.36)		
VD	1.29	(0.34)	75.05	(115.94)		

Table 5: Selected estimates

(b) Store quality

(a) Model parameters

Note: The "Estimate" columns provide point estimates obtained from the indirect inference estimator outlined in Section 6 whereas the "SE" columns report standard errors. I compute standard errors for estimates of the parameters using an analytical expression for the asymptotic variance of indirect-inference estimators; see Online Appendix O.4 for details. I then compute standard errors for the median search costs (in dollars) using the delta method. Each "Med. SC" figure is the median search in dollar terms for a particular retailer among consumers with household income under \$75,000. Additionally, Γ_{fg} is the mean value of γ_i among consumers with initial state h_{i1} given by $h_{ig1} = 1$.

Table 6 reports various descriptive statistics computed on both the estimation sample and on search outcomes simulated from the estimated model to facilitate an assessment of model fit. The table indicates that the model closely fits moments of the estimation sample.

8 Counterfactual analysis

The primary value of this article's model and techniques is in evaluating how search frictions and various forms of seller differentiation shape consumer behaviour and market power in markets with costly search. In this section,

Table 6: Model fit

	Share visiting	Mean $\#$	Share	buying	g from	Share paying	Mean
	one store	of visits	any	1800	VD	> min. price	overpay
Observed	0.82	1.20	0.61	0.36	0.22	0.66	3.95
Baseline	0.84	1.18	0.55	0.34	0.20	0.67	4.13

Notes: the table compares observed and simulated search efforts. "Share paying > min. price" reports the share of purchases occurring at a price above the minimum available price for the consumer's brand whereas "Mean overpay" reports the mean difference between the transaction price and the minimum available price.

I conduct such an evaluation.

8.1 Sources of limited consideration

To understand sources of limited search, I simulate search under counterfactual consumer preferences and assess resulting changes in consumer behaviour. This procedure involves simulating search effort outcomes 50 times for each consumer conditional on prices, prescriptions, and initial states. The counterfactual preference changes include

- (i) Reducing the median search cost from its estimated value to zero.
- (ii) Reducing the state dependence parameter ϕ from its estimated value to zero.
- (iii) Reducing vertical differentiation. This involves setting each retailer f's quality Q_f to $r\hat{Q}_f + (1-r)\bar{Q}$, where \hat{Q}_f is f's estimated quality, \bar{Q} is sales-weighted average quality across retailers, and $r \in [0, 1]$. I reduce r from one to zero.
- (iv) Reducing horizontal differentiation. I do so by setting each consumer's retailer tastes γ_{if} to $r\gamma_{if} + (1 \gamma_{if})\bar{\gamma}_f$, where $\bar{\gamma}_f$ is the unconditional mean of γ_{if} and $r \in [0, 1]$. I reduce r from one to zero.

Figure 2 displays the results. Reductions in search costs, vertical differentiation, and horizontal differentiation all boost consumer consideration. State

dependence plays a smaller role in limiting consideration. Limiting vertical differentiation leads some consumers who previously visited only 1800 to also consider VD, thus raising the average number of visited stores. Conversely, limiting horizontal differentiation leads some consumers who previously visited VD but not 1800 to begin visiting both retailers. This is because visits to VD despite 1800's advantage in terms of quality and search costs require favourable idiosyncratic tastes for VD; limiting these idiosyncratic tastes leads consumers preferring VD in the baseline to begin considering 1800. Reductions of horizontal differentiation eventually reduce consideration because they lead consumers who visited both 1800 and VD to only visit the former.²⁷ Although search costs contribute to limited consideration, only vertical and horizontal differentiation meaningfully influence the extent to which consumers pay above the minimum available price for contacts. Indeed, reducing vertical differentiation lowers mean overpayment whereas reducing horizontal differentiation raises it. The former finding reflects that—as shown by Figure 2c—reducing 1800's quality advantage over VD leads consumers to substitute to the latter store, which generally offers lower prices. Reducing horizontal differentiation has the opposite effect of boosting the mean overpayment. As noted above, consumers often buy from VD rather than 1800 despite the latter store's quality advantage because of idiosyncratic tastes for the former. Weakening these tastes leads VD consumers to substitute to 1800, thus boosting the mean overpayment.

Online Appendix Table O.5 provides results in greater detail for several discrete changes in consumer preferences along with standard errors.

²⁷Online Appendix Figure O.4—which displays changes in the share of consumers visiting VD, 1800, and both retailers as horizontal differentiation is reduced—documents this phenomenon.



Figure 2: Counterfactual search patterns

(a) Mean number of visited retailers

(b) Mean payment over minimum available price (\$)



(c) Difference in market share between 1800 and VD

Notes: the figure plots outcomes of search efforts under counterfactual preference changes described in the main text. The plotted quantities are averages over 5000 simulated search-effort histories (i.e., sequences of distinct search efforts over time) for each consumer in the estimation sample.

8.2 Sources of market power

I assess sources of market power by simulating equilibrium markups under counterfactual consumer preferences using the pricing model of Section 5. Under this model, each store f sets prices p_f to maximize long-run profits

$$\Pi_f(p) = (p_f - mc_f)\sigma_f^L(p)$$

given competitors' prices. In computing pricing equilibria, I use estimates of marginal costs mc_f obtained by solving firms' first-order conditions for profit maximization under observed prices and estimated long-run demand.

The changes in preferences that I consider are discrete versions of the continuous parameter adjustments described in Section 8.1. To lower search costs, I lower each $\bar{\kappa}_f$ to reduce the median search cost for store f by half. To eliminate state dependence, I set $\phi = 0$. To eliminate vertical differentiation, I equalize retailer quality $q_f + \mathbb{E}[\gamma_{if}]$. Last, to eliminate horizontal differentiation, I set $\gamma_{if} = \mathbb{E}[\gamma_{if}|f]$ for each consumer *i* and store *f*. In interpreting the results, note that the approach of incorporating state dependence into a long-run demand system may fail to capture dynamic pricing incentives introduced by state dependence. This concern is most relevant to the counterfactual analysis in which I eliminate state dependence. Additionally, recall that prices are known to consumers before search in the model. This is relevant because, as noted in Section 4.3, the extent of consumer knowledge of prices generally shapes the effect of search frictions on equilibrium prices (see, e.g., Choi et al. 2018). Consumers' uncertainty about prices at the outset of search limits their responsiveness to prices and hence generally inflates markups. This suggests that a model with price uncertainty may predict more negative effects of reducing search frictions on markups.

Table 7 reports effects of counterfactual preference changes on the markups of Acuvue Toric, a popular brand of contacts, in percentage terms. Search frictions do not meaningfully affect retailer market power under the estimated model: reducing search costs does little to change markups. Instead, Table 7 suggests that retailer differentiation drives markups and price dispersion. Eliminating 1800's vertical advantage leads to a 20.4% reduction in its markup, increases in rivals' prices, and an overall reduction in markups. This result implies that vertical differentiation sustains price dispersion. Shaked



Figure 3: Markup changes across brands (medians and IQRs)

Notes: this plot displays the interquartile range (i.e., 25th and 75h percentile) and median of counterfactual markup changes across brands in the estimation sample. It does so separately for 1800 and for VD.

and Sutton (1982) argue that scope for quality differentiation softens price competition by allowing firms to select different quality levels and appeal to market segments with different tastes for quality. This argument seems applicable to contact lens e-commerce based on my results. Horizontal differentiation contributes by far the most to the average markup level, which falls by 54.9% upon its elimination. Figure 3 plots the distribution of markup changes across brands. This figure shows that the results for brands other than Acuvue Toric are similar to those reported in Table 7.

Stone	Low search	No state	No vert.	No horiz.
Store	costs	dependence	diff.	diff.
1800	-1.1 (0.3)	-1.4 (1.1)	-20.4 (3.1)	-47.3(3.6)
WM	5.5(2.0)	1.9(1.1)	22.6(7.0)	-11.0 (36.9)
VD	-0.8 (0.4)	-3.4(1.1)	26.2(8.8)	-73.1 (6.1)
Average	-0.7 (0.2)	-2.0(1.0)	-2.2 (1.7)	-54.9(2.3)

Table 7: Counterfactual markup changes

Note: This table presents estimates of percentage changes in markups for Acuvue Toric under counterfactual consumer preference changes. "Average" provides a sales-weighted average of retailer-specific changes. The standard errors, which appear in parentheses, were computed using a parametric bootstrap with 100 bootstrap draws.

8.3 Sources of the market leader's dominance

The article's techniques also permit a quantification of the extent to which a dominant firm's position in a market owes to an awareness advantage or a quality advantage. The dominance of 1800 in online contact lens sales in 2007–2018 could owe to either sort of advantage. Whereas greater awareness likely reflects 1800's investment in advertising, perceived superior quality likely reflects 1800's investments in logistics and customer service. I quantitatively assess these explanations by simulating search behaviour under the following changes in consumer preferences: (i) equalization of search costs by setting 1800's search cost parameter equal to that of VD, (ii) equalization of quality by changing q_{1800} so that store quality $q_f + \mathbb{E}[\gamma_{if}]$ is equalized across 1800 and VD, and (iii) equalizing both search costs and quality.

Figure 4 displays results. The figure shows that a quality differential rather than an awareness differential primarily explains 1800's sales advantage over VD: equalizing the retailers' search cost distributions reduces the ratio of 1800's to VD's sales from 1.71 to 1.67, whereas equalizing retailer quality reduces this ratio to 0.38. The figure also shows that 1800's quality advantage underlies consumers' choices to buy contact lenses above their minimum available prices—equalizing quality roughly halves mean overpayment—whereas 1800's awareness advantage plays a negligible role in inducing consumers to pay over the minimum available prices for their prescribed brands.

9 Conclusion

This article applied a consumer search model to a panel dataset describing browsing and purchasing in contact lens e-commerce. One contribution of the article is its development of a tractable empirical framework for studying panel sequential search models. This framework exploits a property of the Weitzman (1979) search strategy and, optionally, a convenient set of parametric assumptions to simplify the computation of probabilities of particular



Figure 4: Awareness versus quality differentials

(b) Payment over min. available price

(a) 1800's sales advantage

Notes: this figure provides estimates of (a) the ratio of 1800's to VD's sales and (ii) the mean consumer payment for lenses in excess of the minimum available price for their prescribed brand under counterfactual parameters as described in the main text. The dots provide point estimates whereas the bars provide 95% confidence intervals. I compute the confidence intervals using a parametric bootstrap procedure with 100 replicates.

search outcomes. Another contribution is in drawing substantial conclusions about limited consideration and market power in e-commerce. The analysis suggests that both search costs and seller differentiation explain limited search, but that only the latter accounts for market power in e-commerce.

Bibliography

- Allen, Jason, Robert Clark, and Jean-François Houde. 2014. ""The Effect of Mergers in Search Markets: Evidence from the Canadian Mortgage Industry"." American Economic Review (10): 3365–3396.
- Armstrong, Mark. 2017. "Ordered Consumer Search." Journal of the European Economic Association 15 (5): 989–1024. 10.1093/jeea/jvx017.
- Brown, Zach Y, Mark L Egan, Jihye Jeon, Chuqing Jin, and Alex A Wu. 2023. "Why do index funds have market power? quantifying frictions in the index fund market." Technical report, National Bureau of Economic Research.
- Brynjolfsson, Erik, and Michael D Smith. 2000. "Frictionless com-

merce? A comparison of Internet and conventional retailers." *Management* science 46 (4): 563–585.

- **Chamberlain, Gary.** 1980. "Analysis of covariance with qualitative data." *Review of Economic Studies* 47 (1): 225–238.
- Choi, Michael, Anovia Yifan Dai, and Kyungmin Kim. 2018. "Consumer Search and Price Competition." *Econometrica* 86 (4): 1257–1281.
- Clay, Karen, Ramayya Krishnan, and Eric Wolff. 2001. "Prices and price dispersion on the web: evidence from the online book industry." *Journal of Industrial Economics* 49 (4): 521–539.
- Clemons, Eric K., Il-Horn Hann, and Lorin M. Hitt. 2002. "Price dispersion and differentiation in online travel: an empirical investigation." *Management Science* 48 (4): 534–549.
- Collard-Wexler, Allan. 2013. "Demand fluctuations in the ready-mix concrete industry." *Econometrica* 81 (3): 1003–1037.
- Comscore. 2007-2008. "Web Behavior Database." https://wrds-www. wharton.upenn.edu/pages/about/data-vendors/comscore/, Accessed on December 1, 2018.
- De Los Santos, Babur, Ali Hortaçsu, and Matthijs R. Wildenbeest. 2012. "Testing models of consumer search using data on web browsing and purchasing behavior." *American Economic Review* 102 (6): 2955–2980.
- Dinerstein, Michael, Liran Einav, Jonathan Levin, and Neel Sundaresan. 2018. "Consumer Price Search and Platform Design in Internet Commerce." American Economic Review 108 (7): 1820–1859.
- Dubé, Jean-Pierre, Günter Hitsch, and Peter E. Rossi. 2009. "Do switching costs make markets less competitive?" Journal of Marketing Research 46 (4): 435–445.

- Dubé, Jean-Pierre, Günter Hitsch, and Peter E. Rossi. 2010. "State dependence and alternative explanations for consumer inertia." *RAND Journal of Economics* 41 (3): 417–445.
- **Dubois, Pierre, and Helena Perrone.** 2015. "Price dispersion and informational frictions: Evidence from supermarket purchases."
- Gouriéroux, Christian, Alain Monfort, and Eric Renault. 1993. "Indirect inference." Journal of Applied Econometrics 8 (S1): S85–S118.
- Heckman, James J. 1981. "Heterogeneity and state dependence." In Studies in labor markets, 91–140, University of Chicago Press.
- Hong, Han, and Matthew Shum. 2006. "Using price distributions to estimate search costs." *RAND Journal of Economics* 37 257–275.
- Honka, Elisabeth. 2014. "Quantifying search and switching costs in the US auto insurance industry." *RAND Journal of Economics* 45 (4): 847–884.
- Honka, Elisabeth, and Pradeep Chintagunta. 2017. "Simultaneous or sequential? Search strategies in the U.S. auto insurance industry." *Marketing Science* 36 (1): 21–42.
- Honka, Elisabeth, Ali Hortaçsu, and Matthijs R. Wildenbeest. 2019. "Empirical search and consideration sets." Handbook of the Economics of Marketing 1 193–257.
- Hortaçsu, Ali, and Chad Syverson. 2004. "Product differentiation, search costs, and competition in the mutual fund industry: a case study of S&P 500 index funds." *Quarterly Journal of Economics* 119 (2): 403–456.
- Jolivet, Grégory, and Hélène Turon. 2019. "Consumer search costs and preferences on the internet." *Review of Economic Studies* 86 (3): 1258–1300.
- Kasahara, Hiroyuki, and Katsumi Shimotsu. 2009. "Nonparametric

identification of finite mixture models of dynamic discrete choices." *Econometrica* 77 (1): 135–175.

- Kim, Jun B., Paulo Albuquerque, and Bart J. Bronnenberg. 2010.
 "Online demand under limited consumer search." *Marketing Science* 29 (6): 1001–1023.
- Koulayev, Sergei. 2014. "Search for differentiated products: identification and estimation." *RAND Journal of Economics* 45 (3): 553–575.
- Krasnokutskaya, Elena, and Katja Seim. 2011. "Bid preference programs and participation in highway procurement auctions." American Economic Review 101 (6): 2653–2686.
- Lee, Kwok Hao, and Leon Musolff. 2021. "Entry into two-sided markets shaped by platform-guided search." Unpublished working paper.
- Mehta, Nitin, Surendra Rajiv, and Kannan Srinivasan. 2003. "Price uncertainty and consumer search: a structural model of consideration set formation." *Marketing Science* 22 (1): 58–84.
- Moraga-González, José Luis, Zsolt Sándor, and Matthijs R.
 Wildenbeest. 2023. "Consumer search and prices in the automobile market." *Review of Economics Studies* 90 (3): 1394–1440.
- Moraga-González, José Luis, and Matthijs R. Wildenbeest. 2008. "Maximum likelihood estimation of search costs." *European Economic Review* 52 (5): 820–848.
- Morozov, Ilya, Stephan Seiler, Xiaojing Dong, and Liwen Hou. 2021. "Estimation of preference heterogeneity in markets with costly search." *Marketing Science* 40 (5): 871–899.
- Mundlak, Yair. 1978. "On the pooling of time series and cross section data." *Econometrica* 46 (1): 69–85.

- Pakes, Ariel, Michael Ostrovsky, and Steven Berry. 2007. "Simple estimators for the parameters of discrete dynamic games (with entry/exit examples)." *RAND Journal of Economics* 38 (2): 373–399.
- Saruya, Hiroki, and Michael Sullivan. 2023. "Cross-channel competition and complementarities in US retail." Unpublished working paper.
- Shaked, Avner, and John Sutton. 1982. "Relaxing Price Competition Through Product Differentiation." *Review of Economic Studies* 49 (1): 3–13.
- Sorensen, Alan T. 2000. "Equilibrium price dispersion in retail markets for prescription drugs." *Journal of Political Economy* 108 (4): 833–850.
- Ursu, Raluca M. 2018. "The power of rankings: Quantifying the effect of rankings on online consumer search and purchase decisions." *Marketing Science* 37 (4): 530–552.
- Ursu, Raluca, Stephan Seiler, and Elisabeth Honka. 2023. "The sequential search model: A framework for empirical research."
- Weitzman, Martin L. 1979. "Optimal search for the best alternative." Econometrica 47 (3): 641–654.
- Wooldridge, Jeffrey M. 2005. "Simple solutions to the initial conditions problem in dynamic, nonlinear panel data models with unobserved heterogeneity." *Journal of Applied Econometrics* 20 (1): 39–54.
- Wooldridge, Jeffrey M. 2010. Econometric Analysis of Cross Section and Panel Data. Volume 1. The MIT Press, , 2nd edition.

A Auxiliary statistics of indirect-inference estimator

Table 8 reports the value of each I-I auxiliary statistic both on the estimation sample and on data simulated from the model.

Statistic		$(\hat{\beta}_n)$	Model
Statistic	Value	SE	$(\tilde{\beta}_n(\hat{\theta}))$
Share visiting 1800	0.688	0.014	0.683
Share visiting WM	0.145	0.010	0.149
Share visiting VD	0.360	0.014	0.348
Share buying 1800	0.337	0.014	0.356
Share buying WM	0.024	0.005	0.019
Share buying VD	0.236	0.012	0.200
Share visiting every store	0.013	0.003	0.013
Inertia share	0.846	0.011	0.844
Inertia reg.: indicator for 1800	0.309	0.011	0.363
Inertia reg.: indicator for VD	0.115	0.010	0.137
Inertia reg.: indicator for WM	0.149	0.011	0.222
Inertia reg.: purchased from store last search effort	0.495	0.017	0.401
Inertia reg.: purchased from store two search efforts ago	0.392	0.018	0.400
Role of lagged price: slope for current price	-0.351	0.252	-0.235
Role of lagged price: slope for lagged price	0.023	0.240	0.093
Price sensitivity: slope	-0.155	0.070	-0.174
Cross-visiting: share of 1800 buyers visiting WM	0.116	0.009	0.124
Cross-visiting: share of 1800 buyers visiting VD	0.033	0.005	0.030
Cross-visiting: share of WM buyers visiting 1800	0.308	0.014	0.389
Cross-visiting: share of WM buyers visiting VD	0.128	0.010	0.191
Cross-visiting: share of VD buyers visiting 1800	0.193	0.012	0.177
Cross-visiting: share of VD buyers visiting WM	0.124	0.010	0.123
Dep. of tastes and prices cond. on initial state: slope	-0.302	0.098	-0.359
Price sensitivity heterogeneity	0.045	0.010	0.054

Table 8: Auxiliary model statistics computed on estimation sample

Notes: This note elaborates on Section 6's description of the auxiliary statistics. "Inertia share" is the share of search efforts with the same first-visited retailer as the consumer's previous effort. "Inertia regression" indicates coefficients from a regression of an indicator for whether a search effort included a visit to store f on store indicators and indicators for whether the consumer bought from f in the previous search effort and in the search effort before that. The dataset for this regression includes three observations for each search effort for which t exceeds three, one for each store. "Role of lagged price" includes coefficients from an indicator for whether a search effort ended in a transaction at 1800 on the price at 1800 during search effort t and during the previous search effort. "Price sensitivity" includes coefficients from a regression of an indicator for whether a search effort ended in a transaction at store f on store indicators and the price at f. The regression dataset includes three observations for each effort, one for each store. "Dependence of tastes and prices conditional on initial state" is the slope coefficient from a regression of an indicator for whether a consumer visited store $q \neq f$ on the ratio of the price at the store f for which $h_{if1} = 1$ to the average price of the consumer's brand across stores. I use the prices from the time of the consumer's first-observed purchase. Last, "Price sensitivity heterogeneity" is the slope coefficient from a regression of $(p_{it}^{\text{trans}} - p_{it}^{\min})/p_{it}^{\min}$ on an in-dicator for consumer *i*'s income exceeding \$75,000, where *t* indicates a transaction, p_{it}^{trans} indicates the transaction price, and p_{it}^{\min} indicates the minimum available price. "SE" reports asymptotic standard errors. "Model $(\tilde{\beta}_n(\hat{\theta}))$ " provides the values of the statistics as computed on data simulated from the model under the baseline parameter estimates.