

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

Michael Sullivan

Harvard Business School

July 12, 2023

Click [here](#) for the online appendix.

Abstract

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 assess these explanations with a consumer search and retailer pricing model that I estimate on data describing web browsing for contact lenses. My approach exploits the data's panel structure in estimating the extent of state dependence and taste heterogeneity. I find that seller differentiation is primarily responsible for limited consideration and market power.

*Email address: msullivan@hbs.edu. I thank Steven Berry, Philip Haile, Katja Seim, Mitsuru Igami, and Soheil Ghili as well as participants in Yale's IO Prospectus Workshop, at the Young Economist Symposium 2021, and at the 5th Doctoral Workshop on the Economics of Digitization for helpful feedback. This paper draws on research supported by the Social Sciences and Humanities Research Council.

1 Introduction

The internet facilitates consumer learning about retailers’ product offerings, yet consumers exhibit severely limited consideration in online markets. Additionally, online markets for undifferentiated goods often feature considerable price dispersion, which suggests the presence of market power.¹

What impedes expansive consumer knowledge and cut-throat price competition in e-commerce markets for homogeneous products (e.g., particular book titles or contact lens boxes)? If internet search were costless and sellers’ product offerings undifferentiated, consumers would compare a product’s prices across all retailers and purchase from the lowest-price seller. That online markets for undifferentiated products feature little search and dispersed prices could reflect that search frictions remain significant online — consumers may buy products 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, reputation, 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

¹See, e.g., Clay et al. (2001), Clemons et al. (2002), Moraga-González and Wildenbeest (2008), Koulayev (2014), and Jolivet and Turon (2019).

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

contact lens e-commerce. This setting is attractive for the study of 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, I also analyze books—a category for which this assumption may not hold—and reach similar conclusions as for contacts.

The analysis draws upon panel data of US web browsing and transactions in in 2007–2008. I first document that consumer consideration is severely limited despite significant cross-seller price variation. In 83% of search efforts for contact lenses, the consumer visits only one retailer. Also, the average transaction price for contact lenses is 16% above the minimum price available among major retailers for the consumer’s prescribed brand. These results are similar for other product categories.

To understand limited consideration and retailer market power, I develop a model of sequential consumer search and of retailer price competition. A one-to-one mapping implied by the optimal search strategy of Weitzman (1979) between search effort outcomes and chains of inequalities relating consumer utilities facilitates estimation. In addition, I adapt established techniques for panel-data analysis to my setting to address an initial conditions problem and an endogeneity problem. State dependence and unobserved heterogeneity are separately identified by standard arguments concerning their distinct implications for choice dynamics. The pricing model features retailers who anticipate long-run effects of pricing decisions.

Indirect-inference estimates of the model imply a median search cost of 74 cents. When various forms of seller differentiation are stripped away from the model, the estimated median search cost rises to \$13 — eliminating factors that limit search require search costs to play a larger role in justifying limited consideration. Reliable estimation of search costs thus requires flexible modelling of retailer differentiation.

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 1.17 to 1.36 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 consumers 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. Instead, I find that consumers pay above the minimum available price largely because they value the superior quality of higher-price retailers.

Additionally, seller differentiation shapes equilibrium markups whereas search costs do not significantly contribute to retailer market power. Eliminating horizontal differentiation reduces markups by 48% on average, whereas eliminating the upscale retailer’s vertical advantage reduces its markups by 23%. Prices at rival retailers rise absent vertical differentiation, thus reducing price dispersion. Together, the results suggest that retailer differentiation is responsible for market power and price dispersion in e-commerce.

1.1 Related literature

This article’s primary 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 mar-

kets.³ Some recent studies account for other factors that limit search, including state dependence (Honka 2014) and persistent preference heterogeneity (Morozov et al. 2021). My model incorporates both of these features. My secondary contribution is the development of techniques for estimating a sequential search model. These techniques draw on Weitzman (1979) and Moraga-González et al. (2022).

Some articles in the empirical search literature relating to my own are De Los Santos et al. (2012), Morozov et al. (2021), Koulayev (2014), Jolivet and Turon (2019), Honka (2014), Honka and Chintagunta (2017), Ursu (2018), and Allen et al. (2014).⁴ Sorensen (2000) studies sources of price dispersion among pharmacies, a setting similar to my own, and Dubois and Perrone 2015 study limited search of supermarkets. My article also relates to a literature on inertia in consumer choice (including, e.g., Heckman 1981, Kasahara and Shimotsu 2009), especially Dubé et al. (2009) and Dubé et al. (2010). Last, my work relates to a literature on platform design in e-commerce, including Dinerstein et al. (2018), who study search within eBay, and Lee and Musolf (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. 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

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.

⁵The 2007 and 2008 panels include about 92 000 and 58 000 households, respectively.

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 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 whereas Vision Direct launched in 2004. Contact lens e-commerce was sizeable by 2007; 1800 made net sales of \$125 million in the first half of 2007. Since 2008, many new retailers have entered contact lens e-commerce.

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 (see Table 1). The prices in the time series do not include shipping fees, although 1800 and VD both waived shipping fees for sufficiently large purchases.⁷ Additionally, the price time series do not account for the rebates that manufacturers offered consumers who purchased a sufficient number of contact lens boxes in a single transaction. That these rebates often varied across retailers complicates analysis of contact lens e-commerce. I ignore rebates on the basis that most consumers were ineligible for rebates based on their purchase quantities: in October 2007, rebates were only available to 1800 consumers who purchased at least eight boxes for eight of the ten most popular brands, but only 7.6% of transactions included eight or more

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.

⁷1800, for example, offered free shipping on orders over \$50.

boxes.

I focus on contact lenses but also analyze books, iPods, Playstation 3 (PS3) consoles, and DVDs. These categories contain products with many observed purchases. Online Appendix O.1 describes these categories' data.

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 A details the procedure. In Section 3, I assess robustness to the choice of the maximum number of days before a purchase for a visit to be included in a search effort. 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 first purchase. This reduces the number of transactions from 1956 to 1160.

In the United States, optometrists and ophthalmologists prescribe contact lenses to their patients after administering exams and fittings. 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. Under 15% of consumers in the sample switch brands.

Table 1: Description of transactions in contact lens data

Brand	# trans.	Weeks b/w trans.	
		Mean	Median
Acuvue 2	188	0.56	0.0
Acuvue Advance	145	0.72	1.0
Acuvue Oasys	129	0.81	1.0
Acuvue Advance for Astigmatism	95	1.05	1.0
Biomedics	57	1.79	2.0
Freshlook Colorblends	56	1.89	1.0
Acuvue 2 Colors	51	1.90	1.5
Soflens 66 Toric	48	1.91	1.0
Focus Night & Day	46	2.29	1.0
O2 Optix	46	2.22	2.0
Other brands	474		
Total	1335		

Notes: “# trans.” reports the number of transactions for each brand. The “Weeks b/w trans.” columns describe the duration of time between observed transactions of each brand.

3 Descriptive analysis

3.1 Overview of data

This section describes the data. To begin, Table 1 reports the number of transactions for the best-selling brands in the sample; there are 42 brands altogether. The interquartile range of transaction prices is \$19.99–38.99, and the median number of boxes purchased is two (one for each eye).

Table 2 reports how often consumers in the sample search for and purchase lenses. Consumers make 2.5 search efforts on average, yet some consumers make many more search efforts. The median time between transactions is 14 weeks. Additionally, consumer search exhibits inertia: in 85% of cases, the first store that the consumer visits in a search effort is the same as the first store visited in the consumer’s previous effort.

Table 3a displays the share of contact lens search efforts involving one, two, and three store visits. The “Baseline” column provides results for

Table 2: Repeated purchasing in contact lens panel

	Mean	Quantiles			
		0.25	0.5	0.75	0.95
N. search efforts	2.47	1	2	3	6
N. transactions	1.65	1	1	2	4
# consumers = 793					

search efforts constructed by including visits to 1800 or VD up to 14 days before a purchase using the algorithm described in Appendix A. The “2 days before” column only includes visits made up to two days before a purchase or another visit. Table 3a reveals that limits of consumer search; 83% of search efforts involve a visit to only one store under the baseline data construction. The table also shows that search efforts are insensitive to the choice of parameters used in constructing search efforts. Table 3b provides results for other categories using a five-days-before definition of a search effort. For all categories except PS3s, over 75% of search efforts involve a visit to only one or two stores.

Consumers visit few stores despite the possibility of saving on lenses by visiting and purchasing from other stores. Table 4a 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 the store with the lowest price among visited sites. On average, the consumer pays 7.1% over the minimum available price among visited sites. Table 4b reports analogous results for other categories. In all categories except PS3s, most consumers pay above the minimum available price.

Some retailers may offer superior customer service or shipping, which could

rationalize why consumers purchase from these retailers over lower-price rivals. Table 5 reports for each retailer its number of transactions in the sample and its average price relative to 1800’s price across transactions. The fact that 1800 outsells VD despite charging higher prices suggests that 1800 is more appealing to consumers in non-price dimensions.

Table 3: Share of search efforts by number of visited stores

(a) Contact lenses			(b) Other categories (5 days before)				
# of visits	Share of sessions		# of visits	Share of sessions			
	Baseline	2 days before		Books	iPod	PS3	DVD
1	0.83	0.84	1	0.74	0.40	0.32	0.49
2	0.16	0.15	2	0.22	0.37	0.27	0.34
3	0.01	0.01	3	0.03	0.21	0.33	0.16
			4+	0.00	0.02	0.08	0.00

Notes: see Online Appendix Table O.1 for results for alternative search-effort definitions.

Figure 1 plots coefficients from regressions of (i) overpayment for lenses relative to the minimum price available and (ii) the number of visited stores on various consumer characteristics. The results show that high-income and highly educated consumers purchase contact lens at relatively high prices and that high-income consumers conduct less search; this could reflect that high-income consumers prefer higher-quality retailers charging higher prices, and visit these retailers without searching lower-quality, lower-price sellers. The correlation of income with search behaviour motivates my specification of preferences varying by income in the model.

3.2 Within versus across store search

The fact that consumers require brand-specific prescriptions to purchase contact lenses makes the assumption that search occurs for a fixed product across retailers credible in my setting.⁸ This assumption may not hold

⁸This assumption is credible in other settings, including the textbooks setting of Hong and Shum (2006) when students must buy a particular textbook to satisfy a course requirement.

Table 4: Transactions above minimum available price

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

(b) Other categories				
	Books	iPod	PS3	DVD
Share of transactions above min price	0.70	0.68	0.18	0.51
Average payment over min price (\$)	2.10	6.19	6.73	1.83
Average payment over min price (%)	35.29	4.07	1.38	13.30

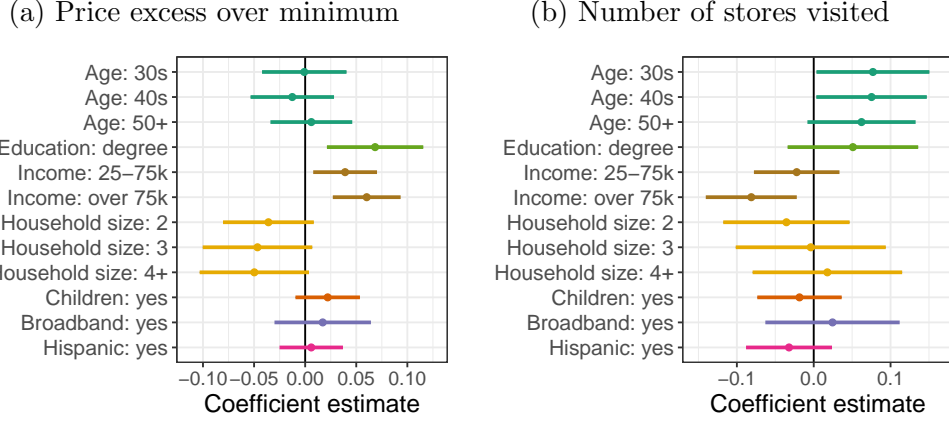
Notes: this table reports the (i) share of transactions made above the minimum available price (MAP); (ii) the average difference of paid price and the MAP, and (iii) the average relative difference of the transaction price over the MAP.

Table 5: Sales and prices by store (contact lenses, 2007–2008)

Store	Transactions	Average relative price
1800contacts.com	849	1.00
visiondirect.com	416	0.85
walmart.com	70	0.94

Note: the average relative price column reports the average ratio of the store's price to 1800's price across transactions in the 2007–2008 sample.

Figure 1: Correlates of search behaviour



Note: Figure 1a plots estimates and 95% confidence intervals from a regression of $(p_t - p_t^{\min})/p_t^{\min}$ on consumer characteristics, where p_t is the transaction price of search effort t and p_t^{\min} is the minimum available price for the consumer's brand during t . Whereas the sample for Figure 1a's regression includes search efforts ending in a transaction, that for Figure 1b's regression includes all search efforts. The R^2 s are 0.053 and 0.022 and the sample sizes are 446 and 675. Estimates of intercepts and race indicators are omitted.

in other product categories. To illustrate, Table 6 describes search within Amazon and Barnes & Noble's online store in search efforts that result in a book purchase. The median number of pages viewed during a visit to Amazon (Barnes & Noble) is 39 (27), and that the median time spent browsing Amazon is about half an hour (16 minutes). Recall also from Table 3 that most consumers do not visit multiple sites before buying a book online. This suggests that consumers search within sites—possibly over distinct products—while conducting little cross-store search. Consumers search much less within contact lenses sites: as reported by Appendix Table O.10, the median number of pages is 12 and the median duration is 6 minutes (pooling across retailers). This is consistent with brand-specific prescriptions that eliminate the rationale for within-store, across-brand search.

Table 6: Within-site search intensity prior to book purchase

(a) Amazon visits					(b) Barnes & Noble visits				
Measure	Mean	Percentile			Measure	Mean	Percentile		
		25 th	50 th	75 th			25 th	50 th	75 th
# pages	102.0	16	39	105	# pages	71.4	10	27	80
Duration	79.9	10	29	79	Duration	48.5	5	16	62

Note: this table reports summary statistics—the mean, 25th percentile, 50th percentile, and 75th percentile—of the number of pages viewed during a visit to an online bookstore and of the duration of time spent (in minutes) browsing the online bookstore.

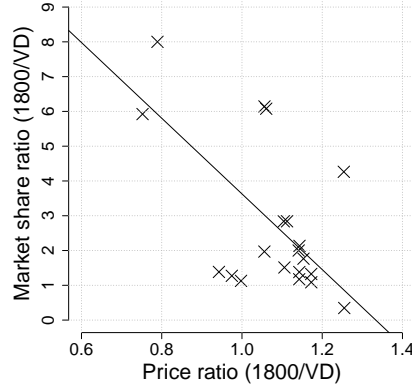
3.3 Prices, browsing, and purchasing

I now turn to the role of prices in directing consumer behaviour. Table 5 shows that 1800 boasted the highest sales despite charging the highest average prices. This could reflect that consumers generally prefer 1800, which leads 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. Figure 2, which plots 1800's sales relative to VD's against its price relative to VD's for the 20 best-selling brands, illustrates this idea: the brands for which 1800 charges especially high prices relative to VD's are those for which 1800's sales are relatively low.

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

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

Figure 2: Prices and intrabrand market shares at 1800 and VD



Note: Each point represents a brand. “Market share ratio (1800/VD)” provides the ratio of transactions at 1800 to those at VD. “Price ratio (1800/VD)” provides the average daily price of a brand at 1800 divided by the analogous quantity for VD. The plot includes the 20 best-selling brands and displays a least-squares line of best fit.

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_{ft} 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.

Table 7 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.⁹ Additionally,

⁹Online Appendix O.6 evaluates the extent to which the positive α estimate reflects cross-brand price differences (i.e., consumers prescribed a brand for which a store charges a relatively high price are less likely to buy from that store on average across time) versus intertemporal price variation, concluding that both sources of price variation are relevant.

Table 7: Descriptive multinomial logit regressions (contact lenses)

Specification 1: $q_{ft} = \bar{q} \quad \forall f, t$			Specification 2: seller/half-year fixed effects		
	Purchase	First visit		Purchase	First visit
α	-0.006 (0.003)	-0.056 (0.010)	α	0.035 (0.004)	0.025 (0.014)
Average elasticity	-0.072 (0.045)	-0.692 (0.086)	Average elasticity	0.449 (0.049)	0.455 (0.111)

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.

the first-visited store responds to prices in a similar way as purchases. This suggests that consumers have some knowledge of prices before search.

4 Model of consumer search

This section outlines the search model. Consumers search for lenses across F online retailers at different occasions in time. Each consumer i has a prescription for a brand j of lenses. The consumer makes search efforts $t \in \{1, \dots, T_i\}$ at exogenously determined times. In each effort, the consumer determines which retailers $f \in \mathcal{F} = \{1, \dots, F\}$ to visit (by, e.g., entering their URLs or making a search-engine query). 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 lenses online (denote $f = 0$). The consumer incurs a search cost κ_{ift} for visiting store f in search effort t . 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

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

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

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 . Additionally, h_{ift} is an indicator for whether the consumer purchased from f in search effort $t-1$. I refer to $h_{it} = \{h_{ift}\}_{f \in \mathcal{F}}$ 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 4.1 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 defined by

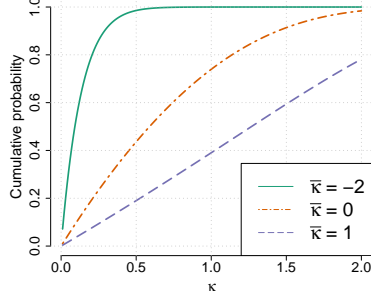
$$\kappa_{ift} = \int_{r_{ift}}^{\infty} (u - r_{ift}) dF_{ift}(u), \quad (4)$$

where $u_{ift} \sim F_{ift}$ conditional on all but ε_{ift} . Note that r_{ift} is the quantity that makes the consumer 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

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

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. (2022) similarly invert equations

Figure 3: Illustration of the search cost distribution function



defining reservation utilities to obtain expressions resembling (4).

There is a convenient parametric distribution of the search costs κ_{ift} that yields tractable choice probabilities for search effort outcomes. Suppose that $\kappa_{ift} \sim F_{\kappa}(\cdot; \bar{\kappa})$ independently of all else, where

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

Then, we can express equation (5) as

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

where the η_{ift} are mutually independent (across i , f , and t) T1EV random variables. The $\bar{\kappa}$ parameter positively relates with both the mean and variance of the F_{κ} distribution. Figure 3 plots F_{κ} for various $\bar{\kappa}$ values.

The distribution above is one of two model features 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. Anal-

ogous 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 of f' and of the outside option in addition to all other reservation utilities. This reasoning is summarized by the following chain of inequalities (wherein I suppress the t subscript):¹⁰

$$r_{if} \geq r_{if'} \geq u_{if} \geq u_{i0} \vee u_{if'} \vee \max_{g \in \mathcal{F} \setminus \{f, f'\}} r_{ig}.$$

Given the distributional assumptions, the search outcome's probability is

$$\begin{aligned} & \frac{e^{\bar{r}_{if}}}{\sum_{g=1}^F e^{\bar{r}_{ig}} + e^{\bar{u}_{i0}} + e^{\bar{u}_{if}} + e^{\bar{u}_{if'}}} \times \frac{e^{\bar{r}_{if'}}}{\sum_{g \neq f} e^{\bar{r}_{ig}} + e^{\bar{u}_{i0}} + e^{\bar{u}_{if}} + e^{\bar{u}_{if'}}} \\ & \times \frac{e^{\bar{u}_{if}}}{\sum_{g \notin \{f, f'\}} e^{\bar{r}_{ig}} + e^{\bar{u}_{i0}} + e^{\bar{u}_{if}} + e^{\bar{u}_{if'}}}, \end{aligned} \quad (7)$$

where $\bar{u}_{ig} = u_{ig} - \varepsilon_{ig}$ and $\bar{r}_{ig} = r_{ig} - \eta_{ig}$. Online Appendix O.2 provides the inequalities corresponding to other outcomes.

The choice probabilities in (7) 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 revealed indirect utilities at each step in search. Last, it would require integration over κ_{ift} and ε_{ift} in order to obtain outcome probabilities. Note that the mapping between chains of inequalities and search effort outcomes reduces the burden of computing choice probabilities even without a parametric assumption on κ_{ift} or the assumption that search costs are iid.¹¹

¹⁰Note that \vee is the maximum operator, i.e. $a \vee b = \max\{a, b\}$.

¹¹Other articles have exploited utility rankings in analyzing search models. Moraga-González et al. (2022) 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.

4.1 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. 2022). It is justified in my context for several reasons. First, regressions from Section 3.3 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.¹² Another reason to assume search over match values is the presence of non-price retailer characteristics that consumers learn through search, especially shipping time. Contact lenses vary not only by brand but also by other prescription parameters; these include base curve, power, sphere, etc.¹³ 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. An alternative approach is to assume search over prices and specify consumer beliefs over prices.¹⁴ 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.

4.2 Justification of sequential search model

Empirical consumer search studies tend to use either sequential or fixed-sample search models. I specify a sequential model for two reasons. First, sequential search better describes cross-store search for contacts. In fixed-

¹²This 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.

¹³Prices do not vary by these parameters.

¹⁴See Mehta et al. (2003), Hong and Shum (2006), Moraga-González and Wildenbeest (2008), and Honka (2014).

sample search, the consumer (i) chooses which stores to visit and (ii) proceeds to visit all of these stores simultaneously before ending search. Due to the technology of web browsing, I expect a consumer to visit one store at a time. If so, the consumer has no reason to commit to visiting every initially selected store or to not search beyond the initially selected stores. Additionally, search order is empirically significant: among consumers who visit both 1800 and VD, those who previously purchased from 1800 visited 1800 before VD 58% (standard error: 3.9%) of the time whereas those who previously purchased from VD visited 1800 before VD 46% (standard error: 3.9%) of the time. This finding contradicts the fixed-sample search model, in which the order of search is irrelevant.

4.3 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 $y_i = \{y_{it}\}_{t=1}^{T_i}$, where y_{it} denotes consumer i 's search/purchase choices in search effort t . Similarly let $p_i = \{p_{it}\}_{t=1}^{T_i}$, where p_{it} denotes the prices of consumer i 's brand at search effort t . Next, let h_{i1} denote consumer i 's initial state, let θ denote an arbitrary parameter vector, and let θ_0 denote the true parameter vector. The model provides conditional probabilities of search effort outcomes that I denote by $\Pr(y_{it}|I_i, p_{it}, h_{it}, \gamma_i; \theta)$. The overall conditional probability of consumer i 's sequence of search efforts

$$\Pr(y_i | I_i, p_i, h_{i1}; \theta) = \int \Pr(y_i | I_i, p_i, h_{i1}, \gamma_i; \theta) dG(\gamma_i | p_i, h_{i1}; \theta),$$

where G is the distribution of γ_i conditional on p_i and h_{i1} .

Integrating over γ_i raises two econometric problems. The first is the standard initial condition problem: the distribution of γ_i conditional on p_i and h_{i1} will depend on h_{i1} because 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 γ_i and prices p_i conditional on h_{i1} . To understand this dependence, suppose that store f sold two brands of contact lenses and that its price for the first brand was high relative to other stores whereas its price for the second brand was relatively low. In that case, consumers with a prescription for the first brand who buy at f require favourable tastes for the store to justify buying from it despite its high price. Similarly, consumers with prescriptions for the second brand may buy from f despite disliking the store to take advantage of 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.7 presents evidence that consumers who previously purchased 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:

$$\gamma_{if} | (p_i, h_{i1}) \sim \begin{cases} N(\lambda \tilde{p}_{jf}, \sigma_\gamma^2), & h_{if1} = 1 \\ N(\Gamma_{fg}, \sigma_\gamma^2), & h_{ig1} = 1 \end{cases}$$

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. First, specifying a parametric distribution of γ_i conditional on the initial state follows Wooldridge (2005).¹⁵ Second, modelling the dependence of γ_i on prices conditional on the initial state follows the correlated random effects (CRE) approach used to address endogeneity in panel data models (Chamberlain 1980, Mundlak 1978, Wooldridge 2010). CRE approaches involve modelling the dependence of unobserved heterogeneity on regressors.

5 Price competition

To analyze market power, I specify a pricing model. The model is static in that each retailer sets a time-invariant price for each brand. The model captures, however, long-run responses of consumer states to prices. An alternative approach is to study Markov perfect equilibria (MPE) of a dynamic game wherein sellers adjust prices in response to changes in payoff-relevant state variables. In my setting, these state variables are the shares of consumers of each (γ_i, α_i) type who previously purchased from each seller. 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. A dynamic pricing model can realistically capture effects of contemporaneous price changes on future sales; the static model, however, captures these effects by accounting for long-run responses of consumer states to prices. Online Appendix O.8 details the dynamic model, which yields results similar to those from the static 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-*

¹⁵As discussed by Wooldridge (2005), the primary alternative is to specify the distribution of the initial state conditional on unobserved heterogeneity, which is far more computationally burdensome than the approach taken here.

run state probabilities $\{\rho_f(p, \gamma_i, \alpha_i)\}_{f=1}^F$, defined as the solutions of

$$\rho_f(p, \gamma_i, \alpha_i) = \sum_g \sigma_{fg}(p, \gamma_i, \alpha_i) \rho_g(p, \gamma_i, \alpha_i) \quad \forall f, \quad (8)$$

where $\sigma_{fg}(p, \gamma_i)$ is the probability with which a consumer with state $h_{igt} = 1$ buys from store f given prices p . The right-hand side of (8) 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 $\rho_g(p, \gamma_i, \alpha_i)$. Thus, condition (8) imposes that the share of type- (γ_i, α_i) consumers in state f does not change after an additional search effort. Letting H denote the unconditional distribution of (γ_i, α_i) , the long-run market share for store f is

$$\sigma_f^L(p) := \int \sum_g \rho_g(p, \gamma_i, \alpha_i) \sigma_{fg}(p, \gamma_i, \alpha_i) dH(\gamma_i, \alpha_i).$$

6 Estimation

6.1 Indirect inference

I estimate the model using an indirect inference (I-I) estimator.¹⁶ This approach involves (i) computing auxiliary statistics $\hat{\beta}_n$ on the sample; (ii) simulating outcomes under a trial parameter value θ using the model; 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 $\tilde{\beta}_n(\hat{\theta})$:

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

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 $\{x_i, h_{i1}\}_i$, outcomes simulated $H = 50$

¹⁶See 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).

times for each panelist. Additionally, $\hat{\Omega}_n$ is a weighting matrix; I use the optimal weighting matrix as discussed in Online Appendix O.5.

I describe the regression coefficients included in $\hat{\beta}_n$ and the parameters that they are included to target in Appendix B. Several of these coefficients are sample averages. To summarize the coefficients:

- (i) *Stores' visit shares*: shares of search efforts with a visit to each store.
- (ii) *Consideration set size*: share of search efforts wherein the consumer visited all stores.
- (iii) *Inertia*: regressions of indicators for whether a consumer visited a store on lagged purchases.
- (iv) *Role of lagged price*: regressions of an indicator for buying from 1800 on the contemporaneous and lagged price at 1800.
- (v) *Price sensitivity*: regression of purchase decisions on prices.
- (vi) *Cross-visiting behaviour*: shares of consumers in various states who visit each store.
- (vii) *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.
- (viii) *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 Table 12 details these statistics and reports their values.

In estimation, I de-mean the prices that enter consumer utilities by the average price across stores conditional on brand and time. Without de-meaning prices, the model would mechanically predict a larger probability

of choosing the outside option for expensive brands.

6.2 Identification

I now informally discuss identification. The endogeneity of price poses a challenge for the identification of price sensitivity. My solution relies on the assumption that retailer quality does not vary across brands, which permits the specification of retailer fixed effects that capture this quality. I could relax this assumption by allowing store quality to vary in other dimensions—e.g., across time or lens characteristics—and specify fixed effects for interactions of stores and, e.g., time periods or characteristics.

The separate identification of state dependence and unobserved heterogeneity γ_{if} is another challenge. Although both state dependence and persistent unobserved tastes 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. Dubé et al. (2010) discuss the use of covariate variation for separate identification. Consider a consumer who buys from store f before f store f raises its price. The consumer responds by switching to purchase from store g , after which f restores its original price. Under extensive state dependence, the consumer is likely to purchase from store g , but if the consumer’s initial purchase from store f stemmed from favourable tastes for f , then we would expect the consumer to switch back to f . These different

implications for switching patterns motivate my “Role of lagged price” I-I auxiliary statistics.

Last, the fact that search costs, state dependence, and persistent store tastes all limit the number of visited stores poses an identification challenge. Separate identification stems from the fact that state dependence and persistent tastes induce choice dynamics that iid search costs do not. Thus, state dependence and persistent tastes are identified by choice dynamics whereas search costs are identified by the extent of search conditional on these former two aspects of preferences. This identification argument requires that search costs are not serially correlated or dependent on past purchases; I leave the relaxation of this assumption to future research.

7 Parameter estimates

Table 8 reports parameter estimates. The “Baseline” panel reports results for the full model, whereas the “Stripped down” panel reports results for a specification without state dependence or persistent heterogeneity. The median search cost under the baseline estimates is only \$0.88, which is low compared to the median transaction price of about \$30. The estimates suggest, however, that taste heterogeneity and state dependence exercise significant influence on consumer decisions. First, the negative estimate of α_1 indicates that high-income consumers (who comprise 37% of the sample) are less price sensitive than their low-income peers.¹⁷ As for state dependence, I find that having previously purchased from a store raises the average consumer valuation of the store by \$4.31, 14% of the mean transaction price. A comparison of the “Baseline” and “Stripped down” results suggests that ruling out state dependence and persistent taste heterogeneity leads to an overstatement of search costs. When these aspects of consumer preferences are ignored, the model requires higher search costs

¹⁷See Online Appendix Table O.12 for elasticity estimates.

to rationalize highly limited search.

Table 9 reports estimates of store quality $q_f + \mathbb{E}[\gamma_{if}]$. 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.

Heterogeneity γ_{if} in tastes for retailers 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.11 and Figure O.3 describe the regression described above in greater detail.

Online Appendix O.9 reports estimates for books e-commerce. The results are qualitatively similar to those for contact lenses — the estimated median search cost is \$0.90, close to the estimate of \$0.74 for lenses.

Table 10 reports various descriptive statistics computed on both the esti-

Table 8: Selected parameter estimates

Parameter	Baseline		Stripped down	
	Estimate	SE	Estimate	SE
q_{1800}	1.061	0.110	1.249	0.160
q_{WM}	-2.469	0.243	-0.646	0.220
q_{VD}	0.185	0.055	0.296	0.161
ϕ	0.468	0.146	-	-
α_0	0.147	0.017	0.069	0.028
α_1	-0.072	0.022	-	-
$\bar{\kappa}$	-1.815	0.097	0.603	0.044
$\Gamma_{1800,VD}$	-3.340	0.251	0.000	0.000
$\Gamma_{VD,1800}$	-5.381	0.424	0.000	0.000
σ_γ^2	1.175	0.038	-	-
λ	3.492	0.801	-	-
Median search cost (utils)	0.11	0.01	0.96	0.03
Median search cost (\$)	0.74	0.13	13.85	5.71

Note: The “Estimate” columns provide point estimates obtained from the indirect inference estimator outlined in Section 6 whereas the “SE” columns report standard errors. Additionally, Γ_{fg} is the mean value of γ_i among consumers with initial state h_{i1} given by $h_{ig1} = 1$.

Table 9: Estimates of store quality

Store f	Mean taste for store f $Q_f = q_f + \mathbb{E}\gamma_{if}$
1800	0.12
WM	-2.96
VD	-3.67

mation sample and on search outcomes simulated from the estimated model to facilitate an assessment of model fit. In general, the model fits moments of the estimation sample well.

8 Counterfactual analysis

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

Table 10: Model fit

	Share visiting one store	Mean # of visits	Share buying from any 1800 VD			Share paying > min. price	Mean overpay
Observed	0.82	1.20	0.61	0.36	0.22	0.66	3.95
Baseline	0.84	1.17	0.73	0.50	0.21	0.71	4.36

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.

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 4 displays the results. Reductions in search costs, vertical differentiation, and horizontal differentiation all boost consumer consideration, with state dependence playing 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. 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. In-

deed, reducing vertical differentiation lowers mean overpayment whereas reducing horizontal differentiation raises it. The former finding reflects that—as shown by Figure 4c—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 — when consumers agree on 1800’s quality advantage over VD, consumers buy from VD and more buy from 1800.

The Online Appendix provides standard errors for the counterfactual estimates. This appendix also includes Table O.14, which provides results in greater detail for several discrete changes in consumer preferences. Additionally, Online Appendix Table O.8 reports results for the books category, which are similar to those for contact lenses. State dependence, however, plays a larger role in limiting search than persistent unobserved tastes in the books category.

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

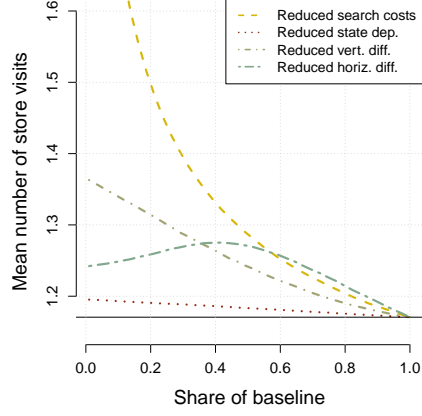
$$\Pi_f(p) = (p_f - mc_f)\sigma_f^L(p)$$

given the prices of its competitors. In practice, I use σ_f^L under the model estimates and estimates of marginal costs mc_f obtained by solving firms’ first-order conditions for profit maximization under observed prices and estimated long-run demand. Throughout this section, I focus on competition in sales of the toric variety of the popular Acuvue label.

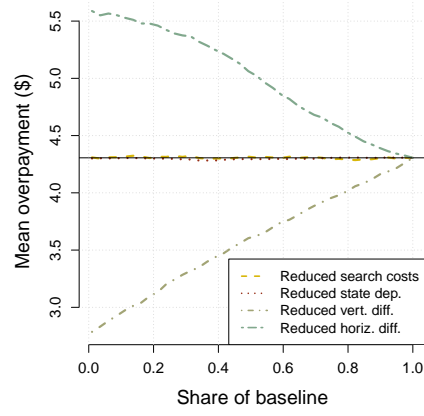
The changes in preferences that I consider are:

- (i) Low search costs: reduce $\bar{\kappa}$ so that the median search cost equals one

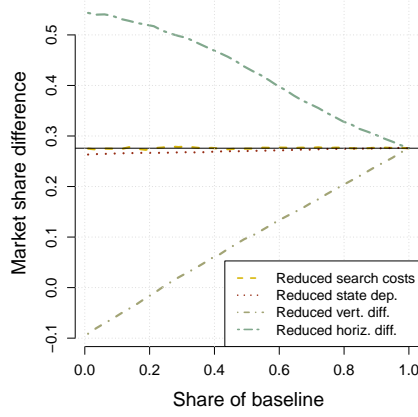
Figure 4: Counterfactual search patterns



(a) Mean number of visited retailers



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



(c) Difference in market share between 1800 and VD

half of the median search cost under the estimated value of $\bar{\kappa}$;

- (ii) No state dependence: set $\phi = 0$;
- (iii) No vertical differentiation: equalize retailer quality $q_f + \mathbb{E}[\gamma_{if}]$ across retailers f at the sales-weighted average retailer quality; and
- (iv) No horizontal differentiation: set $\gamma_{if} = \mathbb{E}[\gamma_{if}|f]$ for each consumer i and each store f .

Table 11 reports effects of counterfactual preference changes in percentage

terms. Note that 1800's relative markup (i.e., price divided by cost) is estimated as 1.87, whereas 1800's sales divided by its costs of goods and services as reported on its Q2 2007 SEC Quarterly Report was 1.69, which is somewhat similar to the estimated relative markup. Reducing search costs does little to change markups, implying that price dispersion for homogeneous goods sold online is not a consequence of search frictions providing sellers with market power. Instead, Table 11 suggests that retailer differentiation drives markups and price dispersion. Eliminating 1800's vertical advantage leads to a 23% reduction in its markup, increases in rivals' prices, and an overall reduction in markups. This result implies that vertical differentiation sustains price dispersion. Shaked and Sutton (1982) 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 48% upon its elimination.

To a small extent, state dependence contributes to retailer market power. The contribution of state dependence to market power has implications for business practices that make switching between stores difficult. A reduction of state dependence could be achieved by, for example, an intermediary service to which the user uploads prescription and delivery information to be shared with other online retailers, thus reducing the hassle of switching. This proposed service resembles e-commerce platforms that provide an interface through which a consumer can interact with many retailers.

Search costs are small and do not meaningfully contribute. As such, easing search using, e.g., comparison tools or transparency regulations, is unlikely to combat market power. Given that searching across contact lens sites is similar to searching the web for other products, the same conclusion may

Table 11: Counterfactual markup changes

Panel A: Point estimates (%)

Store	Low search costs	No state dependence	No vert. diff.	No horiz. diff.
1800	-1.1	-3.4	-23.4	-36.5
WM	5.4	2.0	40.0	-23.7
VD	-1.3	-3.5	47.3	-78.3
Average	-1.0	-3.3	-1.9	-48.1

Panel B: Standard errors

Store	Low search costs	No state dependence	No vert. diff.	No horiz. diff.
1800	0.1	1.1	3.4	4.0
WM	1.3	1.6	6.0	6.0
VD	0.3	1.3	9.7	2.1
Average	0.1	1.1	1.9	2.7

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 were computed using a parametric bootstrap with 100 bootstrap draws.

hold for e-commerce more broadly.

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 search outcomes. The article's primary 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. This result suggests further inquiry into the

sources of seller differentiation as a direction for future research.

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.
- Brynjolfsson, Erik, and Michael D Smith.** 2000. “Frictionless commerce? 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.
- 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 Hortagsu, 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.** 2022. “Consumer search and prices in the automobile market.” *Review of Economics Studies*, Forthcoming.

- 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.
- 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.

APPENDICES

A Construction of search effort panel

In constructing a search effort around a transaction, I include all visits to 1800 or VM in the K days before the transaction and all visits to WM in the $K' \leq K$ days before. In the baseline specification, $K = 14$ and $K' = 2$. I consider alternative values in Section 3. The reason for using a shorter time window for WM is that consumers may visit Walmart for purposes unrelated to contact lenses; a shorter window may exclude such visits. I also construct a search effort for each visit to 1800 or VM that does not result in a transaction. In doing so, I search for visits to retailers within R days (1800 and VD) or R' days (WM) of this visit, and I assign these visits to the search effort of the initial visit. In the baseline specification, $R = 7$ and $R' = 2$. I proceed to add visits that are within R (1800 and VM) or R' (WM) days of visits that have already been added to the search effort, and I continue to iteratively add visits until no more visits are added. In the books, iPods, PS3s, and DVDs categories, consumers are not limited by prescriptions to buy a particular products. This renders untenable the assumption that a visit not resulting in a transaction represents a search effort for a known product. I therefore do not construct search efforts around such visits in these categories.

B Auxiliary statistics of indirect-inference estimator

The following list describes the I-I auxiliary statistics:

- (i) *Stores' visit shares*: the mean of an indicator for whether the consumer visited each store f . These statistics target the q_f and $\bar{\kappa}$ parameters.
- (ii) *Consideration set size*: the mean of an indicator for whether the consumer

visited all available stores. This statistic targets $\bar{\kappa}$.

- (iii) *Inertia*: 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 store. These statistics target ϕ and parameters governing the distribution of γ_i .
- (iv) *Role of lagged price*: coefficients from an indicator for whether a search effort ended in a transaction at 1800 on the price of the consumer's brand at 1800 during the search effort t and during the previous search. These statistics target ϕ and parameters governing the distribution of γ_i .
- (v) *Price sensitivity*: 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. These statistics target α_i and the q_f parameters.
- (vi) *Cross-visiting behaviour*: for each pair of distinct stores (f, g) , the mean across search efforts in which the consumer's state is given by $h_{igt} = 1$ of an indicator for whether the search effort included a visit to f . These statistics target the parameters Γ_{fg} .
- (vii) *Dependence of tastes and prices conditional on initial state*: the slope coefficient from a regression of an indicator for whether a consumer visited store g on the ratio of the price of the consumer's brand 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. I use each store $g \neq f$ in the regressions. This statistic targets λ .
- (viii) *Price sensitivity heterogeneity*: the slope coefficient from a regression of $(p_{it}^{\text{trans}} - p_{it}^{\text{min}})/p_{it}^{\text{min}}$ on an indicator for consumer i 's income exceeding \$75,000.

Table 12: Auxiliary model statistics computed on estimation sample

Statistic	Value	SE
Share visiting 1800	0.688	0.014
Share visiting WM	0.145	0.010
Share visiting VD	0.360	0.014
Share visiting every store	0.013	0.003
Inertia: indicator for 1800	0.309	0.011
Inertia: indicator for VD	0.115	0.010
Inertia: indicator for WM	0.149	0.011
Inertia: purchased from store last search effort	0.495	0.017
Inertia: purchased from store two search efforts ago	0.392	0.018
Role of lagged price: slope for current price	-0.351	0.252
Role of lagged price: slope for lagged price	0.023	0.240
Price sensitivity: indicator for WM	0.181	0.071
Price sensitivity: indicator for VD	0.377	0.064
Price sensitivity: slope	-0.155	0.070
Cross-visiting behaviour: share of 1800 buyers visiting WM	0.116	0.009
Cross-visiting behaviour: share of 1800 buyers visiting VD	0.033	0.005
Cross-visiting behaviour: share of WM buyers visiting 1800	0.308	0.014
Cross-visiting behaviour: share of WM buyers visiting VD	0.128	0.010
Cross-visiting behaviour: share of VD buyers visiting 1800	0.193	0.012
Cross-visiting behaviour: share of VD buyers visiting WM	0.124	0.010
Dep. of tastes and prices cond. on initial state: slope	-0.302	0.098
Price sensitivity heterogeneity	0.045	0.010

Notes: See Section 6 for a description of the auxiliary statistics. “SE” column reports asymptotic standard errors.

Here, t indicates a transaction, p_{it}^{trans} indicates the transaction price, and p_{it}^{min} indicates the minimum available price for the consumer’s brand at the time of the transaction. This statistic targets α_1 .