6-2016

Mixed Pricing in Online Marketplaces

Katja Seim
University of Pennsylvania

Michael Sinkinson
University of Pennsylvania

Follow this and additional works at: https://repository.upenn.edu/marketing_papers

Part of the [Applied Behavior Analysis Commons](#), [Behavioral Economics Commons](#), [Business Administration, Management, and Operations Commons](#), [Cognitive Psychology Commons](#), [Management Sciences and Quantitative Methods Commons](#), [Marketing Commons](#), and the [Sales and Merchandising Commons](#)

Recommended Citation

Seim, K., & Sinkinson, M. (2016). Mixed Pricing in Online Marketplaces. *Quantitative Marketing and Economics*, 14 (2), 129-155. <http://dx.doi.org/10.1007/s11129-016-9168-3>

This paper is posted at ScholarlyCommons. https://repository.upenn.edu/marketing_papers/374
For more information, please contact repository@pobox.upenn.edu.

Mixed Pricing in Online Marketplaces

Abstract

A rich theory literature predicts mixed strategies in posted prices due to standard price discrimination, search frictions, and various other rationales. While typically interpreted as implying occasional sales or price dispersion, online marketplaces enable a firm to truly use randomization as a tool in pricing, and so such behavior should be expected to arise in online settings. We investigate a case of mixed pricing across a large subset of products on a major e-commerce website. We first test for randomizing behavior, and then construct a model of price discrimination that would generate randomization as optimal behavior. We estimate the model and use it to assess pricing effects of a proposed merger in the industry.

Keywords

price discrimination, online markets, merger analysis

Disciplines

Applied Behavior Analysis | Behavioral Economics | Business | Business Administration, Management, and Operations | Cognitive Psychology | Management Sciences and Quantitative Methods | Marketing | Sales and Merchandising

Mixed Pricing in Online Marketplaces

Katja Seim

The Wharton School, University of Pennsylvania and NBER

Michael Sinkinson*

The Wharton School, University of Pennsylvania and NBER

November 2015

Abstract

A rich theory literature predicts mixed strategies in posted prices due to standard price discrimination, search frictions, and various other rationales. While typically interpreted as implying occasional sales or price dispersion, online marketplaces enable a firm to truly use randomization as a tool in pricing, and so such behavior should be expected to arise in online settings. We investigate a case of mixed pricing across a large subset of products on a major e-commerce website. We first test for randomizing behavior, and then construct a model of price discrimination that would generate randomization as optimal behavior. We estimate the model and use it to assess pricing effects of a proposed merger in the industry.

JEL classification: D43, L11, L41, L81

Keywords: Price discrimination, online markets, merger analysis

*The Wharton School, University of Pennsylvania, 3620 Locust Walk, Philadelphia PA 19104 (e-mail: kseim@wharton.upenn.edu; msink@wharton.upenn.edu). We thank the Dean's Research Fund at the Wharton School for financial support. Elizabeth Oppong and Manasvi Ramanujam provided outstanding research assistance. We would like to thank those that provided helpful comments at IIOC, the Econometric Society, and Designing the Digital Economy Conferences, as well as Judy Chevalier, Joe Harrington, JF Houde, Rob Porter, Bruno Strulovici, Jennifer Valentino-Devries, and others.

1 Introduction

A rich theoretical literature establishes equilibria involving mixed strategies in posted prices for firms in various settings, such as when search costs are present or consumers differ in their willingness to pay and mixed strategy pricing serves to price discriminate. These models of mixed pricing are typically motivated by empirical cases of price mixing either over space, or over time: the search literature going back to [Stigler \(1961\)](#) interprets e.g., geographic price dispersion as a realization of mixed strategies, while the sales literature beginning with [Varian \(1980\)](#) views occasional sales as realizations of mixed strategies. However, the advent of online marketplaces has made the idea of mixed pricing strategies feasible in real time: firms can play pricing strategies that had previously only been theorized. This paper describes and analyzes a case of a firm engaging in mixed pricing behavior for a significant subset of products in a \$10B USD online marketplace.

The industry we study is office supplies. Previous academic work ([Ashenfelter et al. 2006](#); [Cabral 2003](#); [Baker 1999](#); [Warren-Boulton and Dalkir 2001](#)) has investigated pricing practices in this retail market, primarily as they relate to concentration in local markets. For the past few decades, there were three nationwide office supply specialty chains in operation: Staples, Office Depot, and OfficeMax. On Sept 4, 1996, Staples and Office Depot announced a merger, although it was subsequently challenged and halted by the FTC in 1997 out of a concern over substantially reduced competition between office supply superstores in various local markets where the firms directly compete against each other, without facing competitive pressure from unrelated rivals.¹ In early 2013, Office Depot and OfficeMax announced a merger, which was permitted in light of declining demand and increased external competition from substitute retail industries like discount stores, warehouse clubs, and online general merchandisers. Finally in 2015, Staples proposed a merger with the combined Office Depot-OfficeMax entity, which is currently under review with a decision expected December 8th, 2015.

A second aspect of pricing by Staples, the largest player, that has been scrutinized in the popular press such as the Wall Street Journal² (WSJ) is its recent surreptitious tailoring of prices on its website based on its customers' locations, in this case zip codes. Accordingly, the website randomizes between showing a potential buyer one of two prices for a given product - a "high" price and a "low" price - with different probabilities in different zip codes. Consumers visiting the website are not directly aware of which price they are seeing, and computer science research indicates it is quite difficult to detect such price discrimination in online markets more broadly ([Mikians et al. 2012](#)). This strategy may inadvertently harm or help consumers who happen to live

¹Federal Trade Commission v. Staples, Inc., 970 F. Supp. 1066 (D.D.C. 1997) grants the FTC's motion for a preliminary injunction.

²"Websites Vary Prices, Deals Based on Users' Information", Jennifer Valentino-Devries, Jeremy Singer-Vine and Ashkan Soltani, Dec 24 2012.

in areas with particular market conditions by altering prices they face. This may improve efficiency when prices are lowered for certain consumers, or may simply reallocate consumer surplus to firm profits if prices rise. The net effect of this is an empirical object that depend on competitors and local market demand conditions.

Such pricing strategies are to be expected going forward in online marketplaces. However, for policy makers, such price discrimination strategies also have implications for merger analysis. As we show below, one particular local market condition that the firm explicitly takes into account is the consumer's distance to the nearest of its rivals' brick-and-mortar outlets. If two competitors merge and reduce the number of offline locations they operate, the distributional effect of this form of price discrimination will thus also change. For example, the combined OfficeMax-Office Depot corporation announced in 2014 that of its 1,900 retail locations, at least 400 would close by the end of 2016, which would then alter prices faced by online consumers in affected markets.³ If Staples were then to acquire the merged rival, the change in competition could have a significant effect on online pricing.

In this paper, we introduce three datasets on pricing in office supplies: first, a "breadth" sample of the high-price display frequency for a single product across all zip codes. Second, a "depth" sample of the same pricing frequency, together with the high and low prices, across a large set of products, but for a small set of zip codes. Finally, an "offline" sample of in-store prices for a small set of products at the three major chains in 17 different cities, for a total of 85 store locations. We document trends in the first dataset, such as how consumers who are close to a rival store location are more likely to see low prices compared to consumers in zip codes that are far from rival chains. In the second dataset, we test directly for randomization behavior, and show that there are no systematic time trends nor serial correlation in the observed prices. We also show that the pricing strategy is statistically identical among the set of products that exhibit randomization. We show in the third dataset that variation in in-store pricing is small and uncorrelated with variation in online pricing.

We then propose a model of price discrimination that rationalizes firm randomization. Our model has "loyal" and "disloyal" consumers, as well as a search parameter among "disloyal" types. We calibrate certain demand parameters and estimate the remaining structural parameters relating to the randomization strategy. Our estimates show that the implied share of loyal customers for Staples.com is increasing in the distance to the nearest rival chain, and that disloyal types exit the market more quickly as the number of establishments in the zip code increases. Our estimates imply small but reasonable distance elasticities.

We then proceed with an analysis of how such online pricing behavior changes after a merger of Staples' rivals. We use data on post-integration store closures of the merged Office Depot-

³"Office Depot, Inc. Announces First Quarter 2014 Results", Office Depot Press Release, May 6 2014.

OfficeMax entity to re-estimate optimal price discrimination for Staples.com, and find that average prices displayed increase by a modest amount compared to pre-merger levels. This can be explained by the fact that rival store closures have the primary effect of removing duplication in the chains' footprints. However, this is important to highlight as merger review in the industry focused on in-store pricing changes from an increase in market power and elasticities of substitution between the three specialty players and unrelated online competitors, but not on within-firm multi-channel pricing.

This paper revisits an extensive theoretic literature on search and sales, as well as growing empirical search and sales literatures. It makes a contribution to the literature on testing mixed strategies ([Walker and Wooders, 2001](#); [Chiappori, Levitt and Groseclose, 2002](#); [Wang, 2009](#)), and also highlights the existence of online price discrimination and its interaction with offline prices. We review the related literature in more detail below.

The paper is organized as follows. Section 2 reviews related literature. Section 3 describes the setting and current merger activity. Section 4 introduces our datasets. Section 5 tests for randomization. Section 6 analyzes the relation between online price mixing and the offline market structure, and rival prices more generally. Section 7 describes and estimates a model of firm pricing behavior. Section 8 concludes.

2 Literature on Mixed Pricing

There is a rich theory literature that explores firms using mixed strategies in equilibrium. We focus on strategies for posted price, recognizing that the auction literature has studied mixed strategies extensively in the context of bids.

A first stream of literature, going back to [Stigler \(1961\)](#), explores price dispersion and the implication for consumer search. [Salop and Stiglitz \(1977\)](#) and [Salop \(1977\)](#) considered how heterogeneity in search costs can lead to price dispersion as a tool for price discrimination. Subsequent work such as [Burdett and Judd \(1983\)](#) and [Stahl \(1989\)](#) modeled consumer search behavior more closely. Recent contributions have revisited competitive aspects ([Janssen and Moraga-González, 2004](#)) as well as informational settings ([Akin and Platt, 2013](#)). There have been a number of important empirical contributions that have used price patterns to identify search costs, such as in [Hong and Shum \(2006\)](#) or [Hortaçsu and Syverson \(2004\)](#), and the relative importance of search in observed prices ([Wildenbeest, 2011](#)).

A second stream of literature on mixed strategies in prices began with [Varian \(1980\)](#) who shows that in competing over “informed” and “uninformed” consumers, firms employ mixed strategies. Subsequent papers such as [Sobel \(1984\)](#) and [Conlisk, Gerstner and Sobel \(1984\)](#) further the intuition that occasionally offering low prices allows a firm to “sweep” the low types of customers from

a market. Later works such as [Raju, Srinivasan and Lal \(1990\)](#), [Rao \(1991\)](#) and [Sinitzyn \(2008\)](#) refined models to showing which types of firms should have low prices more often and how many different prices they needed to charge. A recent contribution, [Chevalier and Kashyap \(2012\)](#), examines both theoretically and empirically price variation across a retailer’s product line, highlighting the interplay between price dispersion and substitution patterns between products. Significant empirical contributions include [Pesendorfer \(2002\)](#) and [Hendel and Nevo \(2013\)](#), which show that intertemporal price discrimination is an effective strategy to differentiate between types in storable goods settings, while [Hendel and Nevo \(2006\)](#) highlights the impact of ignoring such behavior in traditional demand analysis.

In the sales literature, several papers have considered the related decision of uniform (“Every Day Low Pricing”, or EDLP) versus occasional promotion (“HiLo”) pricing, as in [Lal and Rao \(1997\)](#) and [Ho, Tang and Bell \(1998\)](#); [Fassnacht and Hussein \(2013\)](#) offers a review. Other work has explored sale pricing through the lens of reference prices, such as in [Heidhues and Koszegi \(2008\)](#), where the authors show that a monopolist selling to risk-averse consumers can maximize profits by varying between “high” and “low” prices.

Empirical work in mixed strategies has primarily focused on testing for mixed strategies, particularly in sports ([Chiappori, Levitt and Groseclose, 2002](#); [Walker and Wooders, 2001](#)) but also in conduct ([Wang, 2009](#)). These papers test assumptions of randomized play, including the predictive power of covariates and the absence of serial correlation. In addition, they test model predictions of equality of payoffs. We are not aware of any research looking at randomization of posted prices in real time.

We take no stand on which particular thread of literature explains the behavior we document. Instead, we focus on testing for the behavior and documenting its patterns across products and its relation to offline prices. We build a model inspired by the “sweeping” price discrimination mechanism of [Sobel \(1984\)](#) to gain a sense of the magnitude of potential pricing effects of market structure, in the context of the past and proposed mergers.

3 Setting

As of 2013, there were 9,000 office supply specialty stores in the US. Most of these are small, standalone retailers, such as stationary stores, but 42% are owned and operated by one of the three large supply super store chains, Office Depot, OfficeMax and Staples, whose combined estimated market share is upward of 70%. The majority of the firms’ revenue derives from the sale of school and office supplies (45%) and office machines (29%), but they also sell computers, furniture and other supplies. Among them, Staples is the market leader, operating 1,846 outlets in 2013, followed by Office Depot with 1,089 and OfficeMax with 823 outlets based on the chain affiliation of

outlets prior to the companies' merger in late 2013 (see ibisworld.com; firm counts derived from companies' 2013 annual reports.)

Similar to other retail industries, the recession of the late 2000s resulted in sharp declines in demand. The retail chains responded to this, and the additional challenge of increasing external competition from substitute industries like discount stores, warehouse clubs, and online retailers by successfully consolidating. In 2013, Office Depot and OfficeMax proposed and obtained approval to merge; in 2015 Staples and the newly formed entity proposed a further merger that is currently undergoing review.

This follows the earlier proposed merger in 1996 by Staples and Office Depot, a case that the FTC successfully challenged based on concerns and econometric evidence that the merger would substantially reduce competition in the retail sale of office supply superstores in various local markets throughout the country where the firms competed directly against each other. The FTC's case at the time applied a narrow market definition that focused on office supply superstores as a separate market from other office supply specialty stores or general big box retailers. Within local MSA markets, the FTC's evidence suggested that (a) Staples and Office Depot had competed on price before, and that (b) in-store pricing was heavily influenced by the number of other office superstores competing in the local area, where fewer stores led to higher prices. [Ashenfelter et al. \(2006\)](#) present some of the detail behind the FTC's evidence, and [Manuszak and Moul \(2008\)](#) develop a two-stage model of pricing for three representative products, using local market structure as a proxy for competition and instrumenting for the latter nonlinearly using the firms' spatial expansion pattern relative to their headquarters as determinants of equilibrium market structure outcomes. Both show evidence in clear support of the FTC's arguments; we rely on their pricing models to identify demand shifters for the empirical specifications we present below.

The FTC's approval of the later Office Depot and OfficeMax merger concluded that in the intervening years, the market for consumer office supplies had changed significantly, as a result of which the merger of the two office supply superstore chains was unlikely to substantially lessen competition in consumable office supplies sales. The FTC's revised market definition included non-office supply superstores such as Wal-Mart and Target, along with club stores like Costco and Sam's Club, and online retailers of office supplies, most notably Amazon.com, that had grown quickly and significantly since the time of the earlier case and now compete with office supply superstores due to larger office supply product lines. The FTC separately investigated the potential impact of the proposed merger on the sale of consumable office supplies to large regional or national business customers on a contract basis, a market with fewer potential suppliers capable of meeting their large purchasing requirements, but concluded that the merger was unlikely to substantially lessen competition in the contract channel.⁴ The office supply contract market is one

⁴See

https://www.ftc.gov/sites/default/files/documents/closing_letters/office-depot-inc./officemax-

focus of the ongoing investigation of the proposed merger of Staples and the combined Office Depot/OfficeMax entity; the 2013 decision noted that the merged company would still face significant competition from Staples for this business.

4 Data and Descriptive Evidence

Our analysis relies on three data sets. The first one derives from the WSJ’s investigation into Staples’ pricing practices.⁵ As part of a larger investigation into online retailers’ price discrimination practices based for example on a customer’s physical IP address location, the WSJ showcased Staples.com as a typical example: during the period of the WSJ’s study from September to December 2012, the report found that Staples.com’s displayed price depended on the zip code of the computer’s IP address location. Zip codes got either a “high” price or a “low” price with some probability, with high and low prices differing by roughly 10% based on a large set of representative products (these included more than 1,000 different products in 10 selected zip codes - typical examples are a Swingline Stapler, a 12-pack of BIC Rollerball Pens, a 15-pack of Staples-Brand Mailing Tubes, a 24-pack of Intertape Masking Tape, and a SnapSafe-Brand Safe). To demonstrate the pricing practice, the WSJ published on its website a dataset it compiled during its investigation that recorded for each of approximately 42,000 US zip codes the fraction of times Staples.com displayed the high price for the Swingline stapler out of 20 simulated visits to the Staples.com website from a potential customer in that zip code. We denote this data set as the “Breadth” sample in our analysis below, which contains, for each zip code the fraction of times out of the 20 shopping occasions a simulated visitor saw the high price appear. We denote this high-price score as $f \in \{0, 0.05, 0.1, \dots, 0.95, 1\}$. Figure 1 displays the distribution of the frequency with which Staples displays the high price for the stapler in 29,545 standard continental US zip codes. For approximately 7.5% of these zip codes, Staples never displays the high price; similarly, for approximately 5% of zip codes, it always displays the high price. The remaining distribution is bimodal, with approximately 35% of zip codes seeing the high price between 5-15% of the time, while Staples.com displays the high price between 65-75% of the time in around 23% of zip codes.

inc./131101officedepotofficemaxstatement.pdf.

⁵Additional credit should be given to computer science researchers who expose such practices, such as Mikians et al. (2012) which shows significant variation in prices resulting from location, cookie data, and equipment. Mikians et al. (2013) presents results of a fascinating experiment that gathered price data from users with a certain browser toolbar, and found widespread price differences on many popular websites.

Figure 1: Empirical Probability Distribution of High Price

Notes: Figure displays empirical probability distribution of observing the high price for a given product in a given zip code on a given day. Distribution constructed across products and zip codes.

We complement the “Breadth” sample from the WSJ with a similar data set that we assembled subsequently to the WSJ’s reporting for a larger set of products. We used the 2012-2013 comScore Web Behavior data base to identify a set of 268 products that comScore customers purchased from Staples.com immediately prior to our data collection effort. We constructed a web scraping routine that began by loading the Staples.com web page, which created a tracking “cookie”. This cookie contained a field called “zipcode”, which stored a value that we believe was estimated based on the machine’s IP address. We overwrote the zip code field with the targeted zip code and then re-loaded the page and saved the resulting html file. Figure 7 shows the contents of a typical “cookie”, with the zip code field located towards the end. We used this routine to query Staples.com eight times over a period of 4 weeks in April and May of 2013 to record prices for the comScore sample of products from a set of 21 zip codes representing each high-price score from the “Breadth” sample, for a total of over 45,000 queries.⁶ We then parsed the resulting html files using regular expressions in Perl to extract the price that was displayed, as well as whether or not the price was specifically marked as a “sale” price. Among the 268 products, 135 show one price across

⁶Later in 2013, the firm redesigned its website and there is currently no detectable effect of the contents of the zipcode field in the cookie. While it is possible that they have discontinued this pricing strategy, potentially in response to the news coverage it received, it is also plausible that they simply store location data on their servers instead of in cookies, making it impossible to manipulate the latter.

the set of test zip codes, 130 show two prices, 3 show three prices. We drop any scrape results that explicitly declare that the product is offered at a “sale” price. Focusing on the products for which we observe two distinct non-sale prices during the data collection period, we observe the high price being displayed in 35.5% of the 8 price collection rounds. The relative frequency with which the high price is shown in a zip code among our 8 price scrapes is correlated with the WSJ’s distribution in Figure 1, although imperfectly due to the different number of scrapes. Similar to the WSJ, we observe a discount from the high to the low price with two modes - 5% and 10% - with an average price difference of 8% and a median of 8.2%. There are few systematic patterns in which of the 268 products are sold using mixed pricing, as opposed to a single uniform price. Correlating the incidence of price mixing with product attributes - the product’s category (office supplies, furniture, electronics excluding computers, computers, etc.), whether or not the product is a Staples’ store brand, and the price point of the product - suggests that mixing is concentrated among office supply products, but controlling for product category, there are no statistically significant differences between the types of products that have uniform prices versus mixed prices. We denote the resulting sample as the “Depth sample”.

Lastly, we collect information on offline prices across a number of local markets to be able to investigate the extent to which Staples’ online prices relate to its own in-store prices or respond to its rivals’ off-line prices. We hired mystery shoppers from the website TaskRabbit.com, an online marketplace that allows users to outsource small jobs to others in their local area, in November 2014. At the time, TaskRabbit had a presence in 18 cities across the continental US. For each of the cities, we assembled a sample of six stores – two for each office supply super store chain – that differed in Staples.com’s high-price frequency score based on the store’s zip code, but that were sufficiently close to each other to limit driving time by the mystery shoppers. We provided shoppers with a standardized set of five products (Post-Its, a stapler, Scotch tape, pens, and address labels), including pictures and SKUs, and instructed them to record the price for each of the products in each of the stores. Separately, we asked them to record the price for the products on each of the chain’s online shopping platform from their own home zip code. The final data collection effort resulted in price information from 85 stores. The store sample is smaller than the potential sample of 108 stores (18 markets with six stores each) for a number of reasons. First, we were unable to collect price information in Denver, CO, due to bad weather; second, not all three chains operate in all markets, for example, Office Depot did not have presence in Boston, MA prior to its merger with OfficeMax; and third, as the data collection took place shortly after the Office Depot/OfficeMax merger, the chains had already begun to re-optimize their merged store network and closed stores in some of the markets. This resulted both in us sending the mystery shoppers to fewer stores in those markets, and them reporting back store closures among the stores we sent them to for six out of our original set of stores. We have no reason to believe that the store reconfigurations or the

chains’ aggregate market exit patterns are systematically related to Staples.com’s online pricing strategy, as we discuss below. For the average store in the offline pricing sample, denoted as “Offline” below, we observe 4.5 price observations, reflecting that despite our attempts to identify popular, standardized products, not all of the products were available in all stores at the time of the data collection. We also dropped price data reported by one of the mystery shoppers for one of the chains since her reported price information indicated that she had recorded prices for incorrect products. We further dropped price information for products that the mystery shoppers reported to be on a general sale; this occurred for 8.61% of store-price observations. The resulting sample thus consists of 410 price observations from 85 stores in 17 markets.

We add various supplemental data to these three main pricing data sets. First, from ReferenceUSA.com, we construct a panel of each of three chains’ store locations across the continental US as of the end of 2012, 2013, 2014, and mid-2015. We use the latest data primarily to investigate in counterfactual simulations how Staples.com’s online pricing strategy would respond to an alternative offline store network that reflects post-merger reconfigurations. We also use ReferenceUSA to collect information on the store networks of the more distant competitors identified by the FTC’s investigation of the Office Depot/OfficeMax merger: we compile store locations for Walmart, Target, Costco, and Sam’s Club as of the date of our Breadth sample, the end of 2012. We use the detailed longitude and latitude information provided by ReferenceUSA for each store’s location, together with population-weighted centroids for each continental US zip code from the Missouri Census Data Center’s MABLE/Geocorr2K data base to calculate straight-line distance between each zip code and store location. From this master set, we then retain the distance between each zip code and the closest store location for each competitor.

Lastly, we use the Census 2010 to collect demographic information at the zip code level, such as income, wages, and presence of businesses.

4.1 Descriptive Evidence

Using the “Breadth” sample, we can explore patterns in the pricing behavior. Table 1 shows descriptive regressions of the probability of getting a “high” price in a zip code on i) the distance from that zip code to the nearest Staples store location, ii) the distance to the nearest OfficeMax or Office Depot location, and iii) the natural log of the number of establishments taken from the census, as a proxy for local demand conditions and competitiveness. Store locations are likely to be endogenous to both rival locations and local demand, and we see in the first three columns that all three variables are related to the probability of being shown a “high” price. When we add all variables at once, in the final column, we see that the probability of seeing a high price is increasing in the distance to rival stores, and decreasing in the log of local establishments. In unreported

regressions, we find that the distance to other less-direct competitors including Target, Costco, and Sam’s Club is not related to the probability of “high” prices in a zip code.

Table 1: Descriptive Regressions

Dependent Variable: Pr(High Price)						
Model:	(1)	(2)	(3)	(4)	(5)	(6)
Distance to Staples	0.0030*** (0.0004)			0.0004 (0.0003)	0.0019*** (0.0003)	0.0002 (0.0002)
Distance to Nearest Rival		0.0075*** (0.0008)		0.0072*** (0.0008)		0.0066*** (0.0008)
Number of Establishments			−0.0755*** (0.0034)		−0.0622*** (0.0031)	−0.0253*** (0.0046)
N	29,545	29,545	29,122	29,545	29,122	29,122
R_A^2	0.106	0.396	0.169	0.398	0.204	0.413

Notes: zip code-level regressions. Statistical significance at the 10%, 5%, and 1% levels are denoted by *, **, and ***.

In Figure 2, we show the relationship between rival distance and high prices graphically. The figure presents two binned scatterplots of the probability of being shown a “high” price in a given zip code on the distance to the nearest rival store, where rivals are defined as OfficeMax and Office Depot. The first panel is the raw data, while the second is residualized by the log of local establishments. In both we see an increasing relationship between distance to rival and the probability of being shown a “high” price. The first panel seems to indicate that the effect is a step-function, in that only rivals up to a distance of roughly 20 miles are relevant, although the second panel shows a more linear relationship, consistent with the fact that rival locations and local demand are confounded in the first.

Figure 2: Descriptive Analysis of Online Prices

Notes: Figures are “binned scatterplots”, constructed by residualizing the data against all listed controls and adding back means, before plotting means of the data within bins. Regression lines are based on all underlying data, not only bin values.

Table 2: Joint Distribution of zip code distance to closest Staples location and closest rival location

Staples location	Rival location					TOTAL
	< 2 m	2-5 m	5-10 m	10-20 m	> 20 m	
< 2 m	2.25	2.86	1.37	0.75	0.58	7.82
2-5 m	3.28	5.26	2.08	1.55	1.16	13.35
5-10 m	1.55	2.62	3.06	3.22	2.07	12.53
10-20 m	1.46	2.04	1.94	35.63	6.64	47.7
> 20 m	0.57	0.99	1.61	8.78	6.66	18.61
TOTAL	9.11	13.78	10.06	49.93	17.12	100

Notes: Table displays percent of zip codes whose distance to the closest Staples location and closest rival location falls within given distance bands in miles.

As we hope to exploit variation in distances from a zip code to Staples and its rivals, it is important to show that the chains are not perfectly co-located with one another. Table 2 shows that while there is positive correlation in location choices of the three chains, there is still significant variation in their location decisions.

One concern about this exercise is that Staples.com could simply be undertaking some experimentation with price to learn demand. While we cannot conclusively disprove it, we believe this is not the case. First, if it were simple A/B testing to learn demand responses, the level of randomization should not correlate with market characteristics. If instead it were a more recent machine learning algorithm to learn demand, such as a multi-arm bandit “explore-exploit” algorithm, it should be expected to correlate with prior beliefs on correct prices; however it would not be necessary to experiment on the entire country, and we would expect systematic patterns in terms of which products are randomized. Therefore, even though we cannot rule out the possibility of experimentation, it appears less likely. Even if this was indeed a form of machine learning experimentation, most of the contributions of this paper, including documenting how online prices correlate with offline market conditions and the interplay between online and offline pricing remain.

5 Testing for Randomization

Using our “Depth” sample of website scrapes, we can test for whether or not the observed pricing behavior is consistent with randomization. As discussed above, different zip codes have different patterns in the randomization. In Table 3, we regress an indicator variable for seeing the “high” price in a given website price query on different sets of fixed effects. As can be seen in the bottom row of the table, product and date-of-scrape fixed effects have little explanatory power, while zip code fixed effects have a significant amount of power.

For this pricing behavior to be considered random, we must also show that there is no systematic pattern with respect to timing. In Table 4, we add time-related fixed effects to see if they offer any explanatory power. As shown, adding hour-of-day, minute-of-hour, and second-of-minute fixed effects do nothing to explain high prices. Furthermore, tests of joint significance of the time fixed effects fail to reject the null hypothesis, indicating that they collectively offer no explanatory power.

Looking across products, we find the pricing behavior to be strikingly consistent, aside from a small set of 6 outliers.⁷ Among the 122 non-outliers, the mean (median) rate of seeing “high” is 0.334 (0.333) with a 5%-95% range of 0.315 to 0.355. In a simple regression of an indicator for

⁷Three outliers that show high prices far more than others are unusual office supply purchases: store-brand aspirin, acetaminophen and bath tissue. For the remaining three outliers, we have far fewer price scrape observations, since they were put on a uniform sale for a large portion of our time period and we drop all prices marked as sale prices.

Table 3: What Explains a “High” Price?

Dependent Variable: Indicator for High Price						
Fixed Effects for...	(1)	(2)	(3)	(4)	(5)	(6)
Product	X				X	X
Date		X		X		X
Zip code			X	X	X	X
N	18,101	18,101	18,101	18,101	18,101	18,101
Adj- R^2	0.009	0.000	0.925	0.925	0.941	0.941

Notes: Linear probability models for observations from the “Depth” sample for products that display two unique non-sale prices in our data.

Table 4: Timing of “High” Prices

Dependent Variable: Indicator for High Price				
Fixed Effects for...	(1)	(2)	(3)	(4)
Zip code	X	X	X	X
Product	X	X	X	X
Hour of Day	X			X
Minute of Hour		X		X
Second of Minute			X	X
N	18,101	18,101	18,101	18,101
Adj- R^2	0.941	0.941	0.941	0.941
Number of Time Fixed Effects	15	60	60	135
Partial-F of Time Fixed Effects	1.14	1.05	1.27	1.16
P-value of Partial-F	0.3127	0.3653	0.0775	0.1026

Notes: Linear probability models for observations from the “Depth” sample for products that display two unique non-sale prices in our data.

high price on product fixed effects, an F-test of equality of all 122 product fixed effects yields a test statistic of 0.10, and an associated P-value of 1.000, indicating that we fail to reject the null hypothesis that all coefficients are equivalent.

Randomization further requires no serial correlation. Table 5 tests for serial dependence along both the zip code and product panel dimensions, finding no predictive power of the last price shown on the current price being shown. At the bottom of the table are test statistics for panel autocorrelation proposed by Wooldridge (2002), pp 282-283. We fail to find evidence of serial correlation along either panel dimension.

As a final illustration of the randomness achieved by changing the zip code field of the site cookie, Figure 3 shows a full set of forty lags of high prices for the entire Depth sample.

In other research that tests for randomization, the authors often establish that the firm was indifferent among the actions it was randomizing among, a direct implication of the model being tested. In this case, we are unable to conduct such a test, as we do not observe sales figures at the

Table 5: Testing for Serial Correlation

Panel Setting:	Dependent Variable: Indicator for High Price	
	Across zip codes, Within Products	Across Products, Within zip code
Specification:	(1)	(2)
Lagged “High” Price Indicator	−0.0006 (0.0007)	−0.0090 (0.0073)
Controls		
Zip code Fixed Effects	X	X
N	17,418	18,080
Panel Autocorrelation Test Statistic	0.001	0.046
P-value	0.9822	0.8330

Notes: Standard errors are clustered at the product level.

different price levels. Therefore, it is impossible to say that the firm is indifferent between showing the “high” and “low” price. However, in many models discussed in Section 2, a firm would not necessarily be indifferent between the two actions: it is the randomization itself that induces higher overall profits from the overall strategy.

6 Relation to Offline Prices

Using the “Offline” sample introduced above we next investigate whether, and if so how, Staples’ online pricing strategy responds to offline prices, or its competitors’ prices. We rely on the data compiled by mystery shoppers hired on the TaskRabbit platform who recorded prices for five preselected products at two of each of the three chains’ outlets in 17 markets, and simultaneously recorded the online prices for the five products on each competitors’ website when accessing the websites from their home zip code. The resulting sample of 410 price observations from 85 retail stores and 232 observations for online prices reveals the following.

First, there is no variation in Office Depot or OfficeMax’ online prices for any of the five products across zip codes. The Staples’ competitors thus do not appear to be practicing spatial price discrimination on their websites.

Second, we also observe limited variation in offline prices. The average coefficient of variation in the five products’ prices across the 17 markets is 0.065 for Staples, 0.032 for Office Depot, and 0.062 for OfficeMax; there are no systematic patterns across products in the degree of price variation across markets: in Staples’ case, Post-It Notes have the largest coefficient of variation in displayed prices with 0.11; for OfficeMax it is Scotch Tape with a CV of 0.11, and for Office Depot it is Address Labels with a CV of 0.06. Across stores, there is virtually no variation in the price of Pilot pens, which could reflect contracting arrangements with Pilot.

Figure 3: Lagged Autocorrelation in High Prices

Given this limited price variation, it is not surprising that Staples' online price score in the sample store's zip code is virtually uncorrelated with the three chains' chosen in-store price; the correlation coefficient between the online price score and in-store prices is -0.03 for Staples' prices, 0.01 for Office Depot's prices, and -0.16 for OfficeMax's prices. We thus conclude that Staples' online price discrimination strategy is not in response to offline prices, or its competitors' online pricing behavior; the evidence suggests instead that online prices are set independently of offline prices, as well as competitors' pricing strategies, even though online price mixing responds to the proximity of retail outlets to the customer, and thus the convenience with which a consumer could shop at a competitor, rather than through Staples' online channel.

7 Model of Firm Pricing Decisions

A number of models could generate equilibria where a monopolist or oligopolist chooses to randomize among a set of prices systematically across markets. We take no stand on validating a particular assumption taken in the literature; instead we propose a simple model, similar to that of [Sobel \(1984\)](#) with a single good and a dominant firm, and with empirical demand objects instead of stylized reservation prices.

7.1 Model

A firm is selling a good that it can acquire at a constant marginal cost c with no capacity constraint, and is considering different prices to charge: a low price $p^L > c$, or a high price $p^H > p^L$. There are two types of consumers in a market. Loyal (high) types make up fraction α of consumers in any time period, and bargain-hunter (low) types make up $(1 - \alpha)$. Both types of consumers have discrete-choice demand of the Logit form used widely in empirical work (McFadden, 1981). One key difference between the two types is that the bargain-hunters do not consider buying a good at a high price, only at a low price. Furthermore, a loyal customer who chooses to buy zero units in a time period exits the market for the good, while with some probability ρ a bargain-hunter who purchased no units in the same period remains in the market until the next period. We interpret this as the bargain-hunter taking a draw from the distribution of prices charged by other firms in the market when shown a high price, and returning to this firm if the price drawn is greater than p^H , which occurs with probability ρ .

Consider a firm that charges a high price at all times. Demand each period comes solely from high-type customers, so expected profits are

$$\pi^H = M \cdot (p^H - c) \cdot \alpha \cdot s^H(p^H) \quad (1)$$

where M is the market size per time period and $s^H(p^H)$ is the Logit demand function for high types.

A firm could instead charge a low price at all times, earning profits from both segments of consumers:

$$\pi^L = M \cdot (p^L - c) \cdot [\alpha \cdot s^H(p^L) + (1 - \alpha) s^L(p^L)] \quad (2)$$

where $s^L(\cdot)$ is the demand function for low-type consumers.

Now suppose this firm is able to randomize prices, showing a high price with probability f . In a time period where the high price is shown, the firm earns π^H . However, due to the behavior of low types, when a low price is shown, the firm earns additional profit from past low-types who purchased zero after seeing a high price. Expected profits are then

$$\pi^{f,L} = M \cdot (p^L - c) \cdot \left[\alpha \cdot s^H(p^L) + (1 - \alpha) \left(\frac{1}{1 - \rho \cdot f} \right) s^L(p^L) \right] \quad (3)$$

Therefore, a firm that follows a randomization strategy earns total expected profits of

$$\pi = f \cdot \pi^H + (1 - f) \pi^{f,L} \quad (4)$$

Intuitively, the randomization acts as a mechanism to separate types. The more “loyal” types

there are in a market (higher α), the more the firm loses from showing a low price, and so it will put a higher probability on the high price. Similarly, as the probability of market exit of the bargain-hunters decreases (higher ρ), the fewer customers are lost from showing high prices, and so the firm will put a higher probability on high price.

To operationalize this model, we specify the demand of person i of type $k \in \{L, H\}$ as having utility and purchase probability given price p of

$$\begin{aligned} u_{ik} &= \delta_k - \beta_k p + \varepsilon_{ik} \\ s^k(p) &= \frac{\exp(\delta_k - \beta_k p)}{1 + \exp(\delta_k - \beta_k p)} \end{aligned}$$

where ε_{ik} is distribute i.i.d. Type I Extreme Value.

7.1.1 Solving the Model

The firm knows all demand parameters when choosing a pricing strategy. Optimal uniform prices come from optimizing either equation 1 (if serving high types only) or equation 2 (if serving both types).

For a firm choosing a randomizing strategy, the firm optimizes by choosing p^H , p^L , and $f \in (0, 1)$. Optimal choices are given from first-order conditions for each variable in equation 4. The optimal p^H is the same price the firm would choose if it were only serving high types. The optimal choices for p^L and f are inter-dependent, as they interact in equation 3.

7.1.2 Discussion

The mechanism at work is that charging multiple prices separates the two types of consumers, allowing for price discrimination. The firm is able to earn the high profits from high types on occasion, without entirely sacrificing the low types. If there were a single rival charging a uniform price, this firm would be better off charging the low price deterministically after an optimal amount of time. However that cannot be an equilibrium as the rival would then want to “sweep” the low types from the market immediately prior. The fact that the bargain-hunter takes a draw from a distribution of rival firms means the firm can do no better than the randomizing strategy.

The behavior of the low types can be interpreted in several ways. One way is that low type demand is given by a discrete-continuous model of demand, as in [Hanemann \(1984\)](#) or [Dubin and McFadden \(1984\)](#), so that instead of “staying in the market” with some probability after observing a high price, these consumers have zero demand at high prices but demand greater than one unit of the good at low prices. Another interpretation is that a low-type searches for a similar but different product

after observing a high price, such as a substitute good within the same product category, where the firm is employing the same randomization strategy on different goods.

Finally, given that the optimal low price p^L is a function of the rest of the parameters, which vary at the zip code level, it should be the case that the low price varies across zip codes. In practice we do not see this, but at our parameter estimates the range of optimal low prices is small, and so it is reasonable to impose that the low price is constant across zip codes.

7.2 Estimation

Our Breadth data show a single high and low price across all markets for a standard Swingline stapler, with f varying across markets in response to local market conditions. Since we do not observe quantity data, we calibrate the demand parameters β_k and δ_k , $k \in \{L, H\}$, to rationalize the high and low prices we observe, \$14.29 and \$15.79 in this case (calibrated values are in Appendix Table 9). We then want to estimate zip code-level α_z , the share of loyal customers, and ρ_z , the search parameter, as a function of local characteristics. Given our reduced-form analysis of covariates, we choose to model them as follows:

$$\begin{aligned}\alpha_z &= \alpha_0 + \gamma \cdot d_z \\ \rho_z &= \rho_0 + \eta \cdot y_z\end{aligned}$$

where d_z is log of the minimum distance from zip code z to either an OfficeMax or Office Depot location, y_z is the log of the number of establishments in zip code z from the census, and $\theta = \{\alpha_0, \gamma, \rho_0, \eta\}$ is a parameter vector to estimate.

Given a parameter vector θ , we can compute the optimal f_z for zip code z as the profit-maximizing probability of showing high prices.

$$f_z(\alpha_z, \rho_z, \theta) = \operatorname{argmax}_f \pi(f, \theta)$$

We then compare that to the observed high price frequency in the data, f_z^0 , and estimate the parameter vector via minimum distance.

$$\hat{\theta} = \operatorname{argmin}_{\theta} \sum_{z=1}^Z (f_z(\theta) - f_z^0)^2$$

Note that the estimation routine does not weight according to zip code populations or other characteristics; all zip codes are weighted equally as each is a realization of the model for an

individual market.

7.3 Results

The results of the estimation exercise are displayed in Table 6.

Table 6: Estimation Results

Dependent Variable: Indicator for High Price

Parameter	Estimate	Standard Error
α_0	-0.4360	(0.0297)
γ	0.7034	(0.0311)
ρ_0	0.9403	(0.0088)
η	-0.0533	(0.0066)
N	24,855	

Notes: Standard errors are computed via subsampling. Distance to competition and log number of establishments are normalized to have means of one in the sample.

These estimates imply that the share of “loyal” types Staples.com has in a zip code is increasing in the distance to rivals, and the probability of searching again among bargain-hunter types is falling in the number of establishments in the zip code. The figures below show the distribution of values of α_z and ρ_z across zip codes.

Figure 4: Estimated Distributions of Loyal Types Across Zip Codes

Figure 6 shows the joint distribution of α and ρ across zip codes, along with the optimal probability of showing a high price. Intuitively, if there are more loyal types, or if the probability of searching again is high, a higher frequency for the high price is optimal.

Figure 5: Estimated Distributions of Search Parameter Across Zip Codes

The estimates imply that engaging in randomization is strictly more profitable than either only serving the high types by always charging a high price in all markets, or the strategy of serving both types by always charging a low price in all markets. Table 7 shows that profits fall from either alternative strategy, although much more so when the firm only serves the high types in the market. This is consistent with Figure 6 and the fact that most markets have an optimal f between zero and one, but that there is more mass in the lower range of f , and many markets have an optimal f of zero.

Table 7: Profits Under Different Strategies

	Change in Average Profit
Base Case	-
Always charge p^H	(19.88%)
Always charge p^L	(1.71%)

Notes: The table presents simple averages across zipcodes in the change in profit from an alternate strategy relative to the randomization strategy.

7.4 Implications and Discussion

Table 8 shows some basic simulations that illustrate the implications of the model estimates. We first show the effect of increasing rival distances by 20% and increasing number of establishments by 20%; both have modest effects on the effective average price shown in the entire sample, although the distance to rival carries more weight. A 20% increase in distance to the nearest rival

Figure 6: Joint Distribution of Loyal Types and Search Parameter

implies a nearly 32% increase in the share of “loyal” types, while increasing the number of establishments in the zip code by 20% decreases the probability of searching again by a little over 1%. The final column uses data on the change in OfficeMax and Office Depot store locations in 2014 as a result of their merger to determine the effect on “loyal” types in a market for Staples.com. The merger affects the distance to the nearest store for 6.3% of zip codes in our sample. The increase is generally modest as the chain eliminates locations that overlap between the two chains.

Table 8: Model Estimates and Simulations

Scenario:	Base Case	Increase Rival Distance 20%	Increase Number of Establishments 20%	Use 2014 Rival Distances
Average α_z	0.3229	0.4253	0.3229	0.3280
Average ρ_z	0.8870	0.8870	0.8772	0.8870
Average f_z	0.3366	0.4427	0.3236	0.3415
Average displayed price	\$14.79	\$14.95	\$14.77	\$14.81
Average transaction price	\$14.69	\$14.85	\$14.68	\$14.70

Notes: Averages are presented across zip codes and are unweighted. The average displayed price is calculated as $f_z \times p_H + (1 - f_z) \times p_L$, while the average transaction price recognizes that some low types remain in the market for multiple periods, resulting in an average transaction price of $f_z \times \alpha \times p_H + (1 - f_z) \times \left(\alpha + (1 - \alpha) \times \left(\frac{1}{1 - \rho \cdot f_z} \right) \right) p_L$.

While the predicted price increase from the merger is small, it can still imply a significant dollar value given the size of the site and the fact it can be applied to a large selection of products. To put these estimates in context, in 2013, Staples reported \$10.4B in online revenue. A back-of-the-envelope calculation says that if their rivals’ merger allows them to raise prices by even 1% on

roughly half of their product selection, with little demand response, they could see a profit increase of \$50M.

8 Conclusion

This paper makes several contributions to the literature on pricing. First, we test for and document a case of a firm using a mixed pricing strategy for posted prices in an online marketplace. Second, we document the relationship between online and offline pricing for a major product market. Third, we estimate a model for online price discrimination consistent with our data and use it to quantify the magnitudes of the competitive forces on pricing. Our results show that online price discrimination can be an effective tool for firms, and that it is sensitive to offline alternatives. This work revisits large literatures in marketing and economics on price dispersion and price discrimination, and has implications for policy as merger analysis should take into account online price effects from industry consolidation. As pointed out by [Mikians et al. \(2012\)](#), online price discrimination is difficult to detect but widespread, and work such as [Mikians et al. \(2013\)](#) shows the challenge in analyzing such behavior.

References

- Akin, S Nuray and Brennan C. Platt. 2013. “A theory of search with deadlines and uncertain recall.” *Economic Theory* 55(1):101–133.
- Ashenfelter, Orley, David Ashmore, Jonathan B. Baker, Suzanne Gleason and Daniel S. Hosken. 2006. “Empirical Methods in Merger Analysis: Econometric Analysis of Pricing in FTC v. Staples.” *International Journal of the Economics of Business* 13(2).
- Baker, Jonathan B. 1999. “Econometric Analysis in FTC v. Staples.” *Journal of Public Policy & Marketing* 18:11–21.
- Burdett, Kenneth and Kenneth L. Judd. 1983. “Equilibrium Price Dispersion.” 51(4):955–969.
- Cabral, Luis M. B. 2003. “Horizontal mergers with free-entry: why cost efficiencies may be a weak defense and asset sales a poor remedy.” *International Journal of Industrial Organization* 21:607–623.
- Chevalier, Judith A. and Anil K. Kashyap. 2012. “Best Prices.” NBER Working Paper.

- Chiappori, P.-A., S. Levitt and T. Groseclose. 2002. "Testing Mixed-Strategy Equilibria When Players Are Heterogeneous: The Case of Penalty Kicks in Soccer." *The American Economic Review* 92(4):1138–1151.
- Conlisk, John, Eitan Gerstner and Joel Sobel. 1984. "Cyclic Pricing by a Durable Goods Monopolist." *The Quarterly Journal of Economics* 99(3):489–505.
- Dubin, Jeffrey A. and Daniel L. McFadden. 1984. "An Econometric Analysis of Residential Electric Appliance Holdings and Consumption." *Econometrica* 52(2):345–362.
- Fassnacht, Martin and Sabine El Husseini. 2013. "EDLP versus Hi-Lo pricing strategies in retailing—a state of the art article." *Journal of Business Economics* 83(3):259–289.
- Hanemann, W. Michael. 1984. "Discrete/Continuous Models of Consumer Demand." *Econometrica* 52(3):541–561.
- Heidhues, Paul and Botond Koszegi. 2008. "Competition and Price Variation When Consumers Are Loss Averse." *The American Economic Review* 98(4):1245–1268.
- Hendel, Igal and Aviv Nevo. 2006. "Measuring the Implications of Sales and Consumer Inventory Behavior." *Econometrica* 74(6):1637–1673.
- Hendel, Igal and Aviv Nevo. 2013. "Intertemporal Price Discrimination in Storable Goods Markets." *The American Economic Review* 103(7):2722–2751.
- Ho, Teck-Hua, Christopher S. Tang and David R. Bell. 1998. "Rational Shopping Behavior and the Option Value of Variable Pricing." *Management Science* 44(12-part-2):S145–S160.
- Hong, Han and Matthew Shum. 2006. "Using price distributions to estimate search costs." *The RAND Journal of Economics* 37(2):257–275.
- 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.
- Janssen, Maarten C. W. and José Luis Moraga-González. 2004. "Strategic Pricing, Consumer Search and the Number of Firms." *The Review of Economic Studies* 71(4):1089–1118.
- Lal, Rajiv and Ram Rao. 1997. "Supermarket Competition: The Case of Every Day Low Pricing." *Marketing Science* 16(1):60–80.
- Manuszak, Mark D. and Charles C. Moul. 2008. "Prices and Endogenous Market Structure in Office Supply Superstores*." *The Journal of Industrial Economics* 56(1):94–112.

- McFadden, D. 1981. *Structural Analysis of Discrete Data with Econometric Applications*. MIT Press chapter Econometric Models of Probabilistic Choice, pp. 198–272.
- Mikians, Jakub, László Gyarmati, Vijay Erramilli and Nikolaos Laoutar. 2012. Detecting price and search discrimination on the Internet. In *HotNets-XI Proceedings of the 11th ACM Workshop on Hot Topics in Networks*. pp. 79–84.
- Mikians, Jakub, László Gyarmati, Vijay Erramilli and Nikolaos Laoutaris. 2013. Crowd-assisted Search for Price Discrimination in E-Commerce: First results. In *CoNEXT*. pp. 1–6.
- Pesendorfer, Martin. 2002. Retail Sales: A Study of Pricing Behavior in Supermarkets. Technical report Rochester, NY: .
- Raju, Jagmohan S., V. Srinivasan and Rajiv Lal. 1990. “The Effects of Brand Loyalty on Competitive Price Promotional Strategies.” *Management Science* 36(3):276–304.
- Rao, Ram C. 1991. “Pricing and Promotions in Asymmetric Duopolies.” *Marketing Science* 10(2):131–144.
- Salop, Steven. 1977. “The Noisy Monopolist: Imperfect Information, Price Dispersion and Price Discrimination.” *The Review of Economic Studies* 44(3):393–406.
- Salop, Steven and Joseph Stiglitz. 1977. “Bargains and Ripoffs: A Model of Monopolistically Competitive Price Dispersion.” *The Review of Economic Studies* 44(3):493–510.
- Sinitzyn, Maxim. 2008. “Technical Note—Price Promotions in Asymmetric Duopolies with Heterogeneous Consumers.” *Management Science* 54(12):2081–2087.
- Sobel, Joel. 1984. “The Timing of Sales.” *The Review of Economic Studies* 51(3):353–368.
- Stahl, Dale O. 1989. “Oligopolistic Pricing with Sequential Consumer Search.” *The American Economic Review* 79(4):700–712.
- Stigler, George J. 1961. “The Economics of Information.” *The Journal of Political Economy* 69(3):213–225.
- Varian, Hal R. 1980. “A Model of Sales.” *The American Economic Review* 70(4):651–59.
- Walker, Mark and John Wooders. 2001. “Minimax Play at Wimbledon.” *The American Economic Review* 91(5):1521–1538.
- Wang, Zhongmin. 2009. “(Mixed) Strategy in Oligopoly Pricing: Evidence from Gasoline Price Cycles Before and Under a Timing Regulation.” *Journal of Political Economy* 117(6):987–1030.

- Warren-Boulton, Frederick R. and Serdar Dalkir. 2001. "Staples and Office Depot: An Event-Probability Case Study." *Review of Industrial Organization* 19:467–479.
- Wildenbeest, Matthijs R. 2011. "An empirical model of search with vertically differentiated products." *The RAND Journal of Economics* 42(4):729–757.
- Wooldridge, J. M. 2002. *Econometric Analysis of Cross Section and Panel Data*. Cambridge, MA: MIT Press.

Appendix

Table 9: Calibrated Demand Parameters

Parameter	Bargain-Hunter	Loyal
δ	1	3
β	0.275	0.255
c	\$10.46	\$10.46

Notes: Cost value is the lowest available price on Amazon Marketplace for the Swingline Stapler featured in the WSJ article. These values generate optimal prices of \$14.29 for the Bargain-Hunter and \$15.79 for the loyal.

Figure 7: Contents of Sample “Cookie” From Staples.com

```
#LWP-Cookies-1.0
Set-Cookie3: SSID=BgAS-B1GAAAABADxXoFRyMEAAvFegVEBAAAAAAAd7ahR8V6BUQC3v
path="/"; domain=.staples.com; path_spec; expires="2014-05-01 18:29:05Z"; ver-
sion=0
Set-Cookie3: SSPV=E2AAAAAAAahCrTEAAAAAAAIAI8QAAAAAAA;
path="/"; domain=.staples.com; path_spec; expires="2014-05-01 21:56:42Z"; ver-
sion=0
Set-Cookie3: SSRT=Ko-BUQE; path="/"; domain=.staples.com; path_spec;
expires="2014-05-01 21:54:50Z"; version=0
Set-Cookie3: sslb=B; path="/"; domain=.staples.com; path_spec; expires="2013-
05-31 18:34:05Z"; version=0
Set-Cookie3: zipcode=04858; path="/"; domain=.staples.com; path_spec;
expires="2033-04-26 18:29:05Z"; version=0
Set-Cookie3: langPref=%2d1; path="/"; domain=www.staples.com; path_spec;
expires="2033-04-26 18:29:04Z"; version=0
```