NET Institute*

www.NETinst.org

Working Paper #19-07

September 2019

An Empirical Analysis of Minimum Advertised Price Restrictions

Danial Asmat

U.S. Department of Justice, Antitrust Division

Chenyu Yang

University of Maryland

* The Networks, Electronic Commerce, and Telecommunications (“NET”) Institute, <http://www.NETinst.org>, is a non-profit institution devoted to research on network industries, electronic commerce, telecommunications, the Internet, “virtual networks” comprised of computers that share the same technical standard or operating system, and on network issues in general.

An Empirical Analysis of Minimum Advertised Price Restrictions*

Danial Asmat[†] Chenyu Yang[‡]

September 30, 2019

Abstract

Recent theory has examined the competitive effects of minimum advertised price (MAP) restrictions: manufacturer policies that can limit the ability of consumers to search for product prices. In this paper, we empirically study the effect of a major electronics manufacturer’s MAP policy on the e-retail prices for its hardware products. Our approach leverages three types of data: contractual MAP values by product from 2011-2013; daily prices across its largest e-retailers; and the frequency of visits to each e-retailer from a representative household panel. We use a model of search with advertised prices to guide two types of findings. First, descriptive patterns of retailer prices are consistent with the market exhibiting consumer search costs, whereby it is costlier to search when price is below MAP than above MAP. Second, reduced form models imply that MAP diminishes the effect of increased retailer competition on decreased price dispersion, whereby prices are up to 6% more dispersed with MAP than without MAP. This is consistent with a model of inter-retailer price discrimination.

JEL Classification: L41, L81, D83

Keywords: search cost; vertical restraint; advertised price; resale price maintenance; RPM; electronic retail

*We thank Debbie Minehart and Francine Lafontaine for helpful conversations, as well as seminar participants at Michigan Ross, Maryland Smith, the 12th Digital Economics Conference at Toulouse, and the 2019 IIOC. We thank Aaron Franz and Gretchen Sileo for research assistance. We are grateful to the University of Rochester-Simon and the NET Institute (www.NETinst.org) for financial support. All errors are our own. The views expressed are not purported to reflect those of the U.S. Department of Justice.

[†] U.S. Department of Justice, Antitrust Division: danial.asmat@usdoj.gov

[‡] University of Maryland, Department of Economics: chny.yang@gmail.com

1 Introduction

The study of competition in retail markets often features two elements that depart from the simplest conceptions of microeconomic models. At the upstream stage, manufacturers can impose vertical restraints that limit the retailer’s ability to set prices or other types of actions. At the downstream stage, consumers are characterized by holding imperfect information that manifests itself in search costs. Recent work has linked these two concepts by studying vertical policies that influence the *information* that retailers can convey to consumers. Minimum advertised price (MAP) policies, a leading example, can dictate the price level that retailers advertise, but not the price that they ultimately charge.

Although these contracts have existed for decades, policy reports suggest that their role may be expanding—in frequency and importance—with the rise of online shopping. The [European Commission \(2017\)](#) has recently undertaken a large-scale survey of competition in electronic commerce. Its analysis draws a link between increased price transparency and reliance on vertical restraints.¹ Many restraints are informational in nature, including limitations to sell through online marketplaces; use price comparison tools; and advertise online. In Britain, the Competition and Markets Authority (CMA) has commissioned a survey ([Oxera Consulting and Accent, 2016](#)) in part to explore the internet’s impact on the incentives to use vertical restraints, reaching similar conclusions.²

In this article, we empirically study how manufacturer-imposed MAP levels can impact prices set by major online retailers in the U.S. Our approach exploits the availability of changes to the MAP for various products made by Seagate Technology, one of the largest digital storage manufacturers in the world. We proceed in three steps.

First, we use daily online retail prices of Seagate hard disk drives (HDDs), and a consumer panel measuring browsing across retail websites, to document descriptive patterns that point to the existence of search costs tied to retailer advertised price. Second, we specify an intrabrand model of how MAP impacts retailer

¹ “As a reaction to increased price transparency and price competition...[there is] increased recourse to agreements or concerted practices between manufacturers and retailers.” As a mechanism to monitor such agreements, the report also points to the increased adoption of high-speed pricing software among manufacturers and retailers.

²The CMA has also levied fines in two recent decisions concerning MAP policies that target online retailers: [mobility scooters \(2014\)](#) and [commercial refrigerators \(2016\)](#).

prices when there are search costs. The model predicts that under mild conditions, MAP reduces the degree to which increased retail competition cuts downstream prices (relative to the absence of the vertical policy). Third, we test this hypothesis using a reduced form model that compares the prices of products that Seagate selected for MAP to those it did not select for MAP, after matching on observable characteristics. We also evaluate the possibility of different selection mechanisms in the data, and find descriptive evidence that MAP is more likely to be used on more popular products.

Descriptive analysis shows that it is important to consider the role of search frictions in studying the effects of the MAP policy. We find that retailers pricing at or below the MAP post *equal prices* on 58% of days, whereas retailers selling above the MAP post equal prices only 11% of days. This suggests that there are search frictions in the market, whereby observing the transactions price is “costlier” than observing the advertised price through a search aggregator or product page. In addition, the browsing data indicates that the majority of consumers who may have shopped for electronics products on these websites searched a single site rather than multiple sites. This is consistent with consumer browsing habits in other online markets.

Next, we specify a version of the theoretical model of price discrimination under MAP developed by [Asker and Bar-Isaac \(Forthcoming\)](#) (hereafter AB). The model features a monopolist manufacturer selling its product to downstream retailers with a two-part tariff. The market contains three types of consumers, who differ exogenously in their search intensity. Our model departs from AB in only one significant way: we add the presence of loyal customers who shop exclusively at a given retailer. This structure yields pure strategy equilibria with price dispersion, even in the absence of a vertical policy. This pattern is important to model because for each MAP and non-MAP product in our sample, the data show persistent price dispersion across retailers and infrequent price changes.³

We use this model to predict the effect of MAP on product pricing, when retailers compete to varying degrees with each other. Specifically, the model implies that when the correlation between a consumer’s search intensity and product valuation is sufficiently negative, additional browsing overlap between sites (“competition”) tends to reduce the price difference between the sites. We show that this effect can be

³We describe these data patterns more fully in the text.

weaker for products priced under MAP than for products not subject to MAP. This occurs because MAP “equalizes” the advertised price across different retailers while maintaining potentially different transaction prices. Consequently, for the mass of consumers who search based on advertised prices, retailers appear homogeneous. Beginning from an equilibrium that features price dispersion, this effect can maintain price differences across retailers in a manner similar to the inter-retailer price discrimination highlighted by AB.

To test whether the data is consistent with our primary hypothesis, we utilize the presence of Seagate HDDs within the same product line that were not subject to a MAP. We measure the overlap in site visits from the browsing panel over the same time period, which proxies for the degree of competition between any two sites offering a product. Consistent with the hypothesis, we find that whereas an additional percentage point of browsing overlap is associated with a \$0.12-0.17 lower price difference between a non-MAP product-site pair, this difference recedes to zero for product-site pairs under the MAP value. This translates to a price dispersion of up to \$2-4, or 3-6% on the most popular websites selling a typical product. The result holds under a variety of specifications including site fixed effects, and we use these to examine its robustness to different endogeneity threats for selection into MAP.

It is worth emphasizing several facets that are outside the scope of this study. Most broadly, our analysis of MAP is limited to documenting price discrimination effects in the online channel. Studying the impact of price discrimination on welfare would require quantity data. In addition, a total welfare analysis would weigh changes in price against service provision, on- and off-line.⁴ We also limit our analysis to intra-brand, intra-product competition by treating each product as a distinct market. The effectiveness of MAP in raising prices à la minimum resale price maintenance (RPM) should fall with heightened inter-brand competition, but we leave a wider analysis to future work. Our browsing data also lack product-level page visits or purchases. This prevents us from directly measuring how often consumers compare prices across websites for the specific products in our dataset.

⁴Costly services for consumer electronics take place mainly at brick-and-mortar rather than online retailers. The possibility of reduced service provision resulting from the elimination of MAP would therefore occur offline. This reduction would still affect online “showroom” purchases: those that value and use service offline, but ultimately purchase online.

To our knowledge, this is the first empirical study of competition when a manufacturer uses a minimum advertised price policy. Several recent papers have studied the price effects of a vertical policy by using variation in the legality of the policy, including [De los Santos, Kim and Lubensky \(2018\)](#) (MSRP), [De los Santos and Wildenbeest \(2017\)](#) (agency pricing), and [Hunold et al. \(2018\)](#) (MFN clauses). Our setting does not feature variation in the legality of MAP. Instead, we compare the prices of Seagate products selected and not selected for MAP, and examine the viability of different hypotheses for selection effects.⁵

Our study can also be read with reference to the older literature on minimum resale price maintenance (RPM). This is because the inter-retailer price discrimination effect of MAP contrasts with the fixed price floor of minimum RPM, distinguishing their use. Empirical contributions include [Gilligan \(1986\)](#), [Ornstein and Hanssens \(1987\)](#), [Ippolito \(1991\)](#), and [Ippolito and Overstreet \(1996\)](#). [Lafontaine and Slade \(2008\)](#) review the empirical literature on RPM and other vertical restraints.

The remainder of the paper is organized as follows. [Section 2](#) details the Seagate MAP policy in question, as well as other facets of the data and HDD market. [Section 3](#) presents descriptive evidence that suggests the role of search costs in the market. [Section 4](#) presents the model, numerical simulations, and the main hypothesis of interest. [Section 5](#) presents the difference-in-differences model and examines possible mechanisms for MAP selection. [Section 6](#) discusses conclusions and further research ideas on this topic.

2 Seagate MAP Policy and Data Sources

Computer hardware and software products have comprised one of the leading categories for online retail sales in the U.S.⁶ HDDs are the cheapest major type of data storage device for computers. The multibillion dollar market for HDDs approximates a duopoly between the global brands Seagate Technology and Western Digital. In 2017, Seagate reported over \$10 billion of revenue from its HDD business, with over \$3 billion attributable to sales in the Americas.

⁵[Leegin \(2007\)](#) replaced per se illegality of minimum RPM policies with rule of reason analysis at the federal level in the U.S. [MacKay and Smith \(2014\)](#) conduct a cross-industry difference-in-differences analysis that uses this variation.

⁶In 2009, Forrester Research estimated that the percentage of online sales relative to total U.S. sales was about 52% for this category, significantly higher than consumer electronics (14%) or apparel and accessories (9%).

Dating from at least 2009, Seagate imposed a minimum advertised price policy on its U.S. retailers. The contract provides a detailed description of the policy and incentives for compliance. As is typical, Seagate defines the terms and holds unilateral enforcement power. Each reseller is free to charge its sales price independently for all products. For products that are subject to MAP, resellers found to display a price below the specified level forfeit promotional funds that support advertising. Advertising is defined to include:

*Internet advertising such as banner, pop-up, and pop-under ads...
Any “level” of a web site above the “shopping cart”*

Beginning in 2011, Seagate also began publicly announcing the MAP levels of products covered under its policy. Our data includes 22 such lists from 2011 through 2013: an average of more than one per month. The announcements identify the covered products by UPC, which include HDDs as well as other types of data storage devices. Covered products list the MSRP and the MAP. The MAP does not change for each product in each month, but each change is downward in direction. Moreover, each MAP reduction occurs jointly with an MSRP reduction, though not always by the same dollar amount.⁷

To exploit the change in MAP levels, we utilize an archived price dataset from Dynamite Data LLC, a provider of global price and other metrics to e-commerce businesses. Like other e-commerce providers, Dynamite Data used web crawling techniques to search prices daily across hundreds of retailers. It offered a variety of data services to commercial clients, including MAP policy violations and subsequent “cease and desist” measures. To distinguish the advertised from transaction price, its crawler simulated the purchase decision through the shopping cart stage. It marketed these datasets to manufacturer and retail clients.⁸

The dataset contains HDD product prices, scraped daily across the largest U.S. online retail websites, from April 2011- April 2013. It consists only of products sold by *first party* retailers, sites that contract directly with the manufacturer. As such, it excludes prices for products available on marketplaces, which

⁷We are not aware of a model that rationalizes the manufacturer’s decision to coordinate changes in MAP to changes in MSRP, as commonly observed. This behavior may deserve further theoretical inquiry.

⁸Because this dataset only contains transactions prices, and not advertised prices, all references to “price” denote the transactions price unless otherwise stated.

are sold through third party agents via distributors. [Israeli, Anderson and Coughlan \(2016\)](#) show that compliance with MAP is 78%-85% for first party retailers, significantly higher than for third party retailers.

The intersection of the products in the MAP policy lists and price data contains HDDs from Seagate’s flagship *Barracuda* series, marketed toward workstations and high performance PCs, and its *Momentum* series, marketed toward laptops and mainstream desktop systems. [Table 1](#) summarizes these nine products, which include a range of storage capacities. During the sample, each product experienced between one and six different MAP values.

Table 1: Product-Site Descriptive Statistics, 2011-2013

Product	# MAP *	Med P _{rt} †	Med # Ret _t
Momentum LP 250 GB	1	54.99	3
Momentum 320 GB	2	77.47	2
Momentum LP 320 GB	3	63.98	5
Momentum LP 500 GB	3	83.29	5
Momentum LP 1 TB	3	97.65	7
Barracuda 500 GB	3	74.99	6
Barracuda 2 TB	4	117.18	7
Barracuda Green 2 TB	4	118.51	1
Barracuda 3 TB	6	175.77	6

* Number of different MAP values during sample

† Median prices across days

The price includes all discounts that are specific to the product, including sales or rebates. It excludes other elements of the final price to consumers, including shipping and taxes. As is typical with technology products, price generally increases with capacity.

We augment the price data with a consumer browsing panel from the online data provider comScore, Inc. Our comScore panel consists of a subset of the random sample of over 50,000 internet users chosen as part of comScore’s flagship database, for three separate months: May 2011; January 2012; March 2013. Each month contains website visits to the ten largest online retailers identified from the price dataset, ranked by visits in [table B1](#): Amazon, Walmart, Best Buy, Newegg, TigerDirect, Fry’s, Rakuten, Micro Center, CDW, and Insight. It does not contain information for which products were browsed within the website, nor whether a purchase was made. We describe these data further in [section 3.2](#) below.⁹

3 Descriptive Analysis of Retailer Pricing with MAP

3.1 Product-Level Pricing

In this section, we present descriptive statistics of retail prices on MAP products, and consumer browsing habits. These statistics point to the existence of search costs tied to retailer advertised price, and motivate several features of the model to be introduced.

To begin, we focus on a subset of the nine Seagate products subject to the MAP policy. [Figure 1](#) depicts retailer prices for four products of different capacity levels, alongside their respective MAP levels, in black. These products are representative of a pattern of price rigidity: whereas there is dispersion across retailers, the frequency of price changes is relatively low (an average of 12.2 days elapse between a retailer changing its price). Price rigidity among online retailers for Seagate HDDs echoes the findings of [Gorodnichenko, Sheremirov and Talavera \(2018\)](#), who analyze a much broader dataset of online products in a similar time

⁹In our correspondence with comScore, its representatives stated that it does not retain any site-product level data for its consumer sample beyond a two year window. To learn more about how Seagate chooses which products to select for MAP, we also obtained browsing data at the site-product level for the single month of November 2018. We discuss in [section 5](#).

Figure 1: Retail Price and MAP Trends by Product: 2011-2013

period. We interpret the price rigidity as reflecting equilibria of a retailer pricing game with pure, rather than mixed, strategies.¹⁰

The graphs also demonstrate at least two other features. The MAP value is sometimes equal to the lowest retail price: this occurs frequently for the Momentus LP 320. On the other hand, the MAP value can be matched by one retailer yet undercut by others: this occurs for several MAP levels on Barracuda 1 TB and Barracuda 3 TB.

To learn more about retailer pricing behavior with respect to the MAP level, [table 2](#) examines all nine products separately. It estimates the probability that any two retailers tie their prices (i) at or below the

¹⁰Formally, we also conduct a “rank reversal” test for mixed strategy pricing. Define rank reversal as the percentage of days in which, for any two retailers i and j selling the same product, the retailer that priced higher on more total days priced lower on that day. For 214 site pairs selling MAP products during the sample period, the weighted rank reversal statistic has a 25th - 75th percentile range of 0.01 - 0.19.

Table 2: Retail-MAP Price Tie Probabilities, 2011-2013

	$P_{ij,t} \leq P_t^{MAP} \text{ }^\dagger$		$P_{ij,t} > P_t^{MAP} \text{ }^\dagger$		Diff ‡	
Momentum LP 250 GB	0.92	(0.008)	0.13	(0.013)	0.79	[37.6]
Momentum 320 GB *			0.14	(0.016)		
Momentum LP 320 GB	1	(0)	0.09	(0.005)	0.91	[173]
Momentum LP 500 GB	0.31	(0.008)	0.15	(0.007)	0.16	[15]
Momentum LP 1 TB	0.73	(0.007)	0.51	(0.026)	0.22	[8]
Barracuda 500 GB	0.48	(0.007)	0.21	(0.011)	0.27	[21]
Barracuda 2 TB	0.68	(0.007)	0.03	(0.003)	0.65	[85]
Barracuda Green 2 TB *			0	0		
Barracuda 3 TB	0.43	(0.006)	0.08	(0.005)	0.31	[41]

† Sample mean and (standard error) of ties. $P_{ij,t}$ = retail price of sites i and j on day t ; P_t^{MAP} = MAP_t .

‡ Difference in sample means and [t statistic].

* No days on which two retailers post prices at or below MAP.

For each product-category above, take any two retailers $i \neq j$ and define the following two terms.

Let $\mathbb{1}(tie)_{ij,t} = 1$ if $|P_{it} - P_{jt}| < 0.01$, and $T_{ij} = \sum_t \mathbb{1}(P_{it} \neq \cdot) \cap \mathbb{1}(P_{jt} \neq \cdot)$

Then $\overline{tie} = \sum_i \sum_j \frac{\sum_{t \in T_{ij}} \mathbb{1}(tie)_{ij,t}}{T_{ij}}$ and $SE_{\overline{tie}} = \frac{s}{\sqrt{\sum_i \sum_j T_{ij}}}$ for each product-category.

MAP; or (ii) above the MAP. The last column displays the difference and t-statistic between the sample mean of both groups. The conclusion stands out: each product is significantly more likely to feature ties at or below the MAP than above the MAP.

We interpret this finding as consistent with the presence of a mass of consumers who do not incur the search cost to uncover the true transactions price when there is a MAP. To see why, consider the probability that two retailers tie their prices at the MAP level on any given day. If the MAP level is set strictly greater than their marginal costs, and consumers observe the transaction prices on both retailers, then either retailer

would have an incentive to undercut its rival incrementally. Moreover, there would be no basis to expect a higher rate of ties when price is below the MAP threshold.

3.2 Retailers and Search

To further explore the nature of consumer search in this market, it is instructive to examine comScore survey data on browsing habits. By providing session-level identity and duration of time spent on website visits, the panel can shed light on how frequently consumers visit multiple websites to compare prices.

Table 3: Multi-site Visit Probabilities
by Time Period and Website Type

$Pr[s_{pt} \geq 2] : \text{All Websites } ^{\dagger}$			
	$t = \text{day}$	$t = \text{week}$	$t = \text{month}$
μ	0.099	0.189	0.353
SE	(0.0003)	(0.001)	(0.001)
$Pr[s_{pt} \geq 2] : \text{Electronics Websites } ^{\ddagger}$			
μ	0.047	0.078	0.139
SE	(0.0004)	(0.001)	(0.001)

[†] s_{pt} = number of different websites visited by panelist p within time interval t .
Week is calendar week and month is calendar month.

[‡] Excludes Amazon and Walmart.

Define the number of different websites visited by panelist p within a time interval t as s_{pt} . [Table 3](#) displays the mean and standard error of the probability of multi-site visits, $s_{pt} \geq 2$, separately for three time intervals and two different website groups. The probability of visiting at least two sites is evaluated relative to the total number of panelists who visited at least one site in the day (week, month). In the top section, the count is permitted to include any of the 10 sites in the sample. Because the two largest websites

by visits, Amazon and Walmart, do not specialize in computer hardware, the bottom section reestimates the probabilities, counting only visits to sites *other* than those two.¹¹

The table shows that multi-site visit probabilities are low. Taking a one week interval as the proxy for a consumer’s purchase decision, the probability of visiting multiple sites is only less than 20%. The figures are similar if the panel is restricted to high-duration visits (table B2). These figures are in line with those cited by Koulayev (2011), who relates that multi-site visit probabilities in other online shopping studies with search costs have ranged from 20% - 33%. Below, we use this statistic to support a modeling assumption that consumers visit a maximum of two websites.

4 Model of Minimum Advertised Pricing with Search

We now introduce the model of minimum advertised pricing with consumer search frictions. Our approach is based on the price discrimination model set forth by AB. The manufacturer sells a homogeneous product to n exogenously given retailers using a two-part tariff. Retailers set prices to consumers with imperfect information. Consumers differ in search type and valuation. In this type of environment, AB show that advertised price restrictions can facilitate price discrimination between retailers. We focus on the related question of under what conditions MAP can maintain inter-retailer price discrimination as the degree of retailer competition increases.

4.1 Preliminaries

4.1.1 Firms

There is a single manufacturer, M , that sells a homogeneous product to up to n exogenously given retailers. M sells the product at 0 marginal cost using a two-part tariff with per-unit rate w and fixed fee T . Taken together, the terms of the offer (w, T) are assumed to be equal across all n retailers. The fixed fee T can be

¹¹Table B1 shows the percentage and count of website visits across all panelists in the sample.

conceptualized as the promotional funding that the manufacturer guarantees to the retailer; this funding is explicitly noted in Seagate's MAP policy.

Each retailer i chooses two prices (P_i, P_i^a) , where P_i is the transactions price and P_i^a is the advertised price. Following AB, assume the following two conditions:

$$(A1) \ P_i^a \geq P^{MAP} \quad \text{Perfect MAP compliance}$$

$$(A2) \ P_i^a \geq P_i \quad \text{Truthful advertising with no upselling}$$

Under (A1) and (A2), it follows that $P_i^a = \max \{P_i, P^{MAP}\}$. Each retailer advertises either its transactions price or (if it exists) the MAP level, whichever is greater. This precludes the possibility of retailers inducing the consumer to visit with a low advertised price, and then raising the price at the shopping cart stage with extra charges such as shipping or handling.

4.1.2 Consumers

Demand for the product is given as follows. Consumers have unit demand and can take on one of two different valuations, low l and high h . Discrete valuations permit a simple representation of equilibria that feature price dispersion and those that do not.

Consumers are heterogenous in their search *type*. We depart from AB in permitting consumers to take on a variety of search types that can depend on the particular retailer. Specifically, there are three different types of consumers: advertised searchers η , loyal customers δ_i , and price searchers σ_{ij} . The probability of each consumer visiting site i , $Pr[s_i]$, is given below:

$$\left\{ \begin{array}{ll} Pr[s_i|\eta] = \frac{1}{n(A)} \text{ if } i \in A & ; A = \{i : P_i^a = \min_i P_i^a\} \\ Pr[s_i|\delta_i] = 1 & ; Pr[s_k|\delta_i] = 0 \ \forall \ k \neq i \\ Pr[s_i|\sigma_{ij}] = Pr[s_j|\sigma_{ij}] = 1 & ; Pr[s_k|\sigma_{ij}] = 0 \ \forall \ k \neq i, j \end{array} \right.$$

Advertised searchers η visit the retailer that displays the lowest *advertised* price among all n retailers selling the product. If there is a tie, then these consumers pick one retailer from among the set of minima with equal probability. Loyal customers δ_i visit only retailer i , ignoring all others. Price searchers σ_{ij} visit two exogenously given retailers i and j . In sum, the market consists of only these three types of consumers: each consumer visits at least one retailer, and no consumer visits more than two retailers.

If $s_i = 1$, then the consumer purchases the product only if his valuation, v , is at least as high as the *transactions* price P_i . Assume that each consumer has $v \geq l$. The consumer of type η or δ_i purchases from the single retailer he visits if $v \geq P_i$, and does not make a purchase otherwise. Type σ_{ij} purchases from P_i if $v \geq P_i$ and $P_i < P_j$; if $P_i = P_j$ then it buys from i or j with equal probability.

Finally, there is an assumed relationship between the consumer's search *type* and *valuation*. Taking $0 \leq \lambda \leq 1$, denote the probability that the consumer has low valuation with λ^S , λ^A , and λ^L for (respectively) price searchers σ_{ij} , advertised searchers η , and loyal customers δ_i . The consumer is willing to pay the high price with the probability $1 - \lambda^{type}$, and the willingness to pay probability is not contingent on the particular retailer visited.

4.1.3 Timing and Information

The timing of the game is identical to AB. First, M chooses (w, T) and whether to impose a MAP or not. Second, each of the n retailers accepts or rejects (w, T) .¹² Third, retailers set (P_i, P_i^a) , where P_i^a equals $\max\{P_i, P^{MAP}\}$. Fourth, consumers visit the relevant retailer(s) based on (P_i, P_i^a) . Fifth, consumers make purchases and firms realize profits.

The manufacturer M is assumed to have complete information on the three types of consumers $\{\eta, \delta_i, \sigma_{ij}\}$, the two types of valuations $\{l, h\}$, and the willingness to pay parameters $\{\lambda^S, \lambda^A, \lambda^L\}$. Retailer i is assumed to know the fraction of its own potential visitors δ_i and σ_{ij} , as well as ad searchers η and the remaining parameters. These informational assumptions are a simple way to conceptualize the idea that M will set

¹²If a retailer rejects the offer, then it pays neither w nor T . If MAP is imposed, all retailers who accept the offer are assumed to comply with the policy perfectly, and all retailers receive the promotional funds.

the two-part tariff to extract most of the profits from its retailers. Consequently, the relevant equilibrium concept is subgame perfect Nash equilibrium (SPE), and M solves the game via backward induction.

4.1.4 Solving for Equilibrium

M solves the following constrained optimization problem:

$$\begin{aligned} \underset{w, T}{\text{maximize}} \quad & \Pi^M = w \cdot \left(\sum_{i=1}^n q_i \right) + \sum_{i=1}^n \mathbb{1}(q_i > 0) \cdot T \\ \text{subject to} \quad & \Pi_i \geq 0 \text{ and } q_i \geq 0 \end{aligned}$$

There are 2^n possible combinations of contracting relationships. Given a per-unit rate w , retailer price vectors must constitute a Nash equilibrium:

$$\Pi_i(P_i; P_{-i}, w) \geq \Pi_i(P'_i; P_{-i}, w) \quad \forall P'_i \neq P_i \text{ and } \forall i$$

M sets the optimal fixed fee $T^*(w) = \operatorname{argmin}\{\Pi_i \mid q_i > 0\}$. If there are multiple Nash equilibria for a given w , then we assume that retailers play the one with highest $T^*(w)$, i.e. the highest manufacturer profit.

The subgame perfect equilibrium is the value of $(w^*, T^*(w^*))$ that generates the maximum Π^M for all elements in the set $\mathcal{W} = \{\frac{w}{100} \leq h : w \in \mathbb{N}^+\}$. We focus on pure strategy equilibria, based on the infrequency of price changes in the empirical data.

In [appendix A.1](#), we characterize the SPE of this game when $n = 2$ retailers. The SPE takes on a discontinuous (piecewise) function depending on the level of h relative to l and the other parameters. We show that when $\delta_1 \neq \delta_2$, i.e. there is retailer heterogeneity, the SPE in pure strategies is “generally unique”: if there is a continuous distribution over the parameter space, the probability that multiple manufacturer choices of $(w^*, T^*(w^*))$ will yield the same manufacturer profit and downstream retail prices is zero.

4.2 Numerical Simulations

It is important to convey the important elements of the model before generating the primary hypothesis of interest. In this section, we numerically simulate how the distribution of search types shapes equilibrium price levels, with and without MAP. To be concrete, define the set $\mathcal{S} = \{0.05, 0.10, \dots, 0.95\}$. We take the x-y grid to be all pairs $(\sum_i \delta_i, \sum_{i,j} \sigma_{ij})$ in the set $\mathcal{T} = \{x, y \in \mathcal{S} : x + y \leq 1\}$. To control the level of heterogeneity across firms, we draw each of the individual parameters from a truncated normal distribution:

$$\delta_i \sim TN \left(\frac{\sum_i \delta_i}{n}, \left(\frac{\sum_i \delta_i}{n^\alpha} \right)^2 \right)$$

$$\sigma_{ij} \sim TN \left(\frac{\sum_{i,j} \sigma_{ij}}{n}, \left(\frac{\sum_{i,j} \sigma_{ij}}{n^\alpha} \right)^2 \right)$$

The mean is proportional to the number of firms, and the variance decreases for higher values of the scaling parameter α . Finally, for each x-y pair, we take the dependent variable z to be the mean over n retailers and d draws: $\sum_{d=1}^D \sum_{i=1}^n \frac{z_{id}}{nD}$.

Figure 2: Mean Retailer Price Levels, $n = 5$ and $d = 20$ draws, low variance

Consider a Monte Carlo simulation over the retailer price vector, and take the exogenous parameters $l = 5$, $h = 6$, $\lambda^S = 1$, and $\lambda^A = \lambda^L = 0$. These parameters are chosen to illustrate a range of price equilibria across the parameter space, and thereby highlight some differences between the MAP regime and default regime. [Figure 2](#) depicts the results of this simulation when $n = 5$ retailers, $\alpha = 1$ (low variance), and $d = 20$ draws for each pair. It conducts the simulation separately for the regime in which M does not impose MAP (a), and when M sets $\text{MAP} = h$ (b).

In both panels (a) and (b), the mean retailer price tends to decrease as the mass of price searchers $\sum \sigma_{ij}$ increases. This is easily visible because of the assumed perfectly negative correlation between search type and valuation, but would also be present when the correlation is weaker.¹³

Consequently, the simulation shows two distinct ways in which MAP can change retail price equilibria. The first is when there are relatively few price searchers, and the non-MAP equilibrium is uniformly high price. For the mass of price searchers between roughly 10% and 20%, MAP can result in a *lower* average price. This is an example of the “price discrimination” effect illustrated in Proposition 1 of AB. Relative to the non-MAP scenario, the MAP scenario permits high-priced retailers to retain an even fraction of the mass of advertised searchers η , who would otherwise visit only the low-priced retailers. This results in a new subgame perfect equilibrium featuring price dispersion rather than price uniformity.

The second way that MAP can change price equilibria is seen by observing the rest of the parameter space, i.e. the region where price searchers are greater than 20%. In this region, price dispersion already exists *without M resorting to MAP*.¹⁴ By comparison to panel (b), the effect of imposing a MAP when prices feature dispersion is generally to shift the distribution of prices higher, i.e. to other dispersed equilibria with more retailers pricing high. This effect is made possible through retailer heterogeneity. Moreover, because price dispersion is common in MAP and non-MAP products within our dataset, this effect is the more important one to consider for our empirical analysis.

¹³As λ^A or λ^L increases, then more equilibria display uniformly low prices.

¹⁴Although it is not strictly implicit from the figure, the underlying data in this region show that less than 10% of the constituent draws feature a uniform price equilibrium.

4.3 Marginal Effect of Competition on Prices

Next, we consider the comparative static of an increase in the share of searchers σ_{ij} on the difference $|P_i - P_j|$, with and without MAP. This proposition assumes $n = 2$ retailers and $\delta_1 > \delta_2$: retailer 1 attracts more loyal customers than retailer 2.

Proposition 1. *For a given parameter vector, assume that the SPE displays price dispersion:*

$P_1 = h$; $P_2 = l$. Under the conditions above, if $h \leq \frac{l}{1-\lambda^S}$:

- (i) *The difference $|P_1 - P_2|$ is weakly decreasing in σ_{12} ;*
- (ii) *There exists a $\hat{\sigma}_{12}$ such that $\forall \sigma_{12} > \hat{\sigma}_{12}$, $P_1 = P_2 = l$;*
- (iii) *Moreover, $\hat{\sigma}_{jk}$ is strictly greater when the manufacturer has imposed MAP than when it has not.*

Proof. See [appendix A.2](#). □

Figure 3: Change in SPE as Overlap Between 1 and 2 Increases

The sufficient condition holds that the fraction of price searchers σ who have high valuation, $(1 - \lambda^S)$, cannot be too large. This condition rules out the possibility that an increase in σ induces the manufacturer to contract to induce an (h, h) equilibrium. It captures the intuition that as consumers search more intensively across websites, in order to lead to a decrease in price dispersion and in overall price levels, consumers with higher search intensity must have sufficiently low valuations.

Figure 3 graphically illustrates the mechanism behind the proposition with an example from one numerical simulation. Panel (a) depicts the manufacturer’s optimization decision in the absence of MAP. On the x-axis, the mass of price searchers σ_{12} is increased from 0 to 0.6. On the y-axis, M ’s optimal wholesale price w^* is depicted in a dotted line. The graph shows that when σ_{12} becomes sufficiently large, M profitably switches to a lower w^* because it backward induces that a lower wholesale price will yield a uniformly low price equilibrium (l, l) . This equilibrium raises M ’s profits by equalizing the fraction of searchers σ_{12} between retailers 1 and 2, permitting a higher fixed fee T^* .

Panel (b) depicts the equilibrium prices of retailers 1 and 2 as σ_{12} increases. The blue line is constant because $P_2 = l$ in all simulations. The red solid line corresponds to P_1 in the non-MAP scenario, which switches to low price at $\sigma_{12} = 0.36$, as implied by panel (a). Finally, the red dotted line corresponds to P_1 when $\text{MAP} = h$. In contrast to the non-MAP case, the SPE does not yield a low price until $\sigma_{12} = 0.48 > 0.36$. By allocating advertised customers evenly across firms, MAP allows the manufacturer to maintain dispersed prices in the interval $[0.36, 0.48]$.

5 Evidence of MAP effect on Retail Prices

In this section, we present empirical evidence of the effect of MAP on prices across websites as the degree of competition changes. To develop this relationship, we consider a dataset of Seagate products that includes those selected for MAP and those unselected for MAP. We use this dataset to estimate a difference-in-differences model that compares the price difference across the same pairs of websites, for products subject to the policy and not. We also discuss potentially confounding factors, and the sign of possible endogeneity biases.

5.1 Regression Evidence

In addition to the products subject to MAP, Seagate sold 10 others in the Momentum and Barracuda lines that are tracked in the price data. These HDDs possess similar combinations of capacity, speed (RPM), and

form factor to those considered thus far. They differ in that Seagate did not select them to take part in the MAP policy.

To proceed, we use a dataset that is at the website pair-product-day level. We group each observation into a “treatment” or “control” group based on the level of prices and MAP status. Treated observations are pair-product-days for which both sites price under the MAP level. Control observations are from products that are not subject to MAP. The effect of competition across websites is proxied by the matrix Σ of site overlap probabilities. There is one matrix for each year of the sample, based on the single month of that year in which comScore data was purchased (e.g. January 2012, see [fig. B1](#)).

After assembling the dataset, we match observations between the treatment and control group. The primary attribute that determines the price of an HDD is the capacity level. To match observations between both sets of groups, we retain only those observations that have at least one treated and one control observation within the same capacity group on the same day. Indexing a website pair-product-day observation by xt , we estimate the following regression:

$$y_{xt} = \beta_0 + \beta_1 \text{overlap}_{xt} + \beta_2 \mathbb{1}(uMAP)_{xt} + \beta_3 \text{overlap}_{xt} \times \mathbb{1}(uMAP)_{xt} \\ + Cap_x + RPM_x + Form_x + \delta_t + \nu_x + \epsilon_{xt}$$

The parameters of interest are the interaction between the percentage overlap of panelists, *overlap*, and the indicator for both prices in the website pair occurring below the MAP value, $\mathbb{1}(uMAP)$. Define Cap_i , RPM_i , and $Form_i$ as fixed effects for the capacity, RPM speed, and form factor of product i respectively; δ_t as month-year fixed effects; and ν_x as website-pair fixed effects.

[Table 4](#) displays the results from this regression. The constant represents the average dollar difference between the same product at two different websites for the excluded capacity group, 500 GB, when the site overlap is zero and the product is not subject to MAP. The capacity fixed effects show that the dollar difference tends to increase for higher capacity products, which are more expensive. The coefficient on *overlap* in specifications (1) and (2) implies a \$0.12-0.17 lower difference for each percentage point of additional

Table 4: Effect of Website Overlap on Price Differences, Nov 2011-April 2013

	(1)		(2)		(3)		(4)	
	Zero and Non-zero pairs included						Non-zero pairs only	
500 GB								
1000 GB	-0.631	(0.467)	0.416	(0.445)	0.458	(0.449)	1.625	(0.371)
2000 GB	1.434	(0.560)	1.843	(0.574)	1.891	(0.490)	4.255	(0.701)
3000 GB	6.700	(1.002)	7.234	(0.770)	7.058	(0.876)	9.250	(0.812)
overlap	-0.117	(0.054)	-0.171	(0.054)	-0.755	(0.767)	-0.566	(0.929)
1.uMAP	-3.277	(0.690)	-3.438	(0.587)	-3.291	(0.572)	-1.896	(0.742)
1.uMAP $\times$ overlap	0.138	(0.071)	0.191	(0.068)	0.115	(0.044)	0.091	(0.037)
Constant	17.289	(5.505)	20.607	(5.594)	19.288	(5.217)	20.476	(5.830)
Observations	29,463		29,463		29,463		21,047	
R-squared	0.217		0.273		0.220		0.271	
Month $\times$ year FE	Y		Y		Y		Y	
RPM and Form Factor FE	Y		Y		Y		Y	
Retailer FE	N		Y		N		N	
Retailer Pair FE	N		N		Y		Y	

Dependent variable equals $|P_{x1,t} - P_{x2,t}|$ for any pair of sites $(x1, x2)$ and days t . Coefficient and (standard error).

Standard errors clustered at website pair-product level. Sample includes only Seagate products in Momentus and Barracuda lines. Begins from the first MAP policy effective date, November 27 2011, through April 18 2013. Price outliers dropped beyond the 95th percentile of website pair-product-day observations, for each capacity group.

website overlap. To gain a sense of its magnitude, consider the site overlaps in table B1. Specification (2) implies that the average price difference of a 500 GB HDD, sold without an MAP, between any two sites with negligible overlap is \$20.61. The average price difference between Amazon and Best Buy on the same product is \$20.61 - $(13.5 \times 0.17) = \$18.30$; between Amazon and Wal-Mart is \$16.43.

With this in mind, the interpretation of the interacted coefficients is straightforward. In both specifications (1) and (2), adding up the coefficients for *overlap* and *uMAP $\times$ overlap* sums to approximately zero. To continue with the example above, this implies that the average price difference *on a product sold under the MAP value* is no lower between Amazon and Best Buy than between any two sites with negligible overlap.

The average price of a 500 GB MAP product sold less than the MAP value during the sample period was \$72. This means that without MAP, the prices of these products may have been lower by up to about \$4, or 1-6 percentage points, on the three most popular sites (Amazon, Wal-Mart, Best Buy).

One type of endogeneity threat to interpreting this regression is non-stationarity in wholesale prices. The structure of our theoretical model assumes that wholesale price w^* is equal across all sites. This would be invalid if, for example, large sites (those with more overlap) obtain better terms and hence price closer together than smaller sites.

To deal with this possibility, specifications (3) and (4) include fixed effects for each website pair; (4) is limited to daily pairs that have a non-zero price difference. The independent variable *overlap* is identified using purely time-series variation. This variation comes from the evolution of the overlap matrix in table B1 across the three months (one for each year) for which we have data. Consequently, the estimation of the *overlap* coefficient becomes very noisy. Nonetheless, the coefficient of interest remains similar in magnitude. This provides some evidence that fixed differences in wholesale costs are not driving this result.¹⁵

5.2 Additional Evidence on MAP Selection

The regression model above does not account for how Seagate selected products to be designated for MAP. Selection effects are important to consider for two reasons. They may present endogeneity threats to interpreting the price discrimination effects discussed above. Moreover, understanding the selection process can yield more insight about the strategic decision to use MAP from the manufacturer’s perspective.

One selection hypothesis is illustrated by simulating the theoretical model in appendix fig. 5. This figure shows that Seagate should be more likely to use MAP on products that feature a lower mass of price searchers than others. This is more likely to hold the more negative is the correlation between search intensity and valuation. Because MAP tends to provide retailers with a greater incentive to price high, price searchers with low valuation are more likely to exit the market when MAP is used. This reduces the total surplus in the market.

¹⁵Clustering at the website pair-product level allows arbitrary autocorrelation across days, within each website pair-product. The magnitude of the estimated variance reflects that prices are sticky: websites did not change prices frequently.

If the comScore sample overlap probabilities overstate the true incidence of multi-site shoppers for MAP products, then the coefficient of interest is biased upward. It was not readily possible to test this hypothesis within the 2011-2013 sample period, because comScore did not retain page-level archives from this period at the time of our inquiry.

To gain further insight, we studied Seagate’s use of the MAP policy in a more recent time period. We obtained November 2018 data on (i) site-product level browsing, and (ii) corresponding Seagate MAP contracts. Beginning from a PC panel of over 50,000 participants, this supplementary comScore dataset tracked each page visit to the 10 websites above. We narrowed this dataset to panelists who visited Seagate products: in November 2018, 254 panelists visited 75 different Seagate products that were subject to MAP, and 28 panelists visited 21 different non-MAP Seagate products.¹⁶

92.5% of the 310 panelists who visited Seagate products visited a single website in that month. Panelists from the MAP and non-MAP groups visited respective averages of 1.06 and 1.03 different websites (p-value=0.53), and 1.76 and 1.75 different Seagate products (p-value=0.98) during the month. Although the power of the test is limited by the small number of panelists who visited non-MAP products, the test does not suggest that panelists who browsed Seagate MAP products in November 2018 were less likely to search intensively.

A second selection hypothesis is that demand is higher for products selected to take part in MAP. If there is some enforcement cost to engage in MAP, Seagate may maximize its profit by “paying” to enforce MAP only on its more popular products.¹⁷ Table 5 summarizes the differences in panelist visits, product ratings, and product reviews between MAP and non-MAP products. Consistent with higher demand, several of these measures suggest that MAP products were more popular than non-MAP products.¹⁸

¹⁶An additional 18 panelists visited both MAP and non-MAP products during this month. They are excluded from this analysis.

¹⁷Another interpretation relates to offline advertising such as newspaper inserts or coupons. Because it is costly to purchase space to advertise each product, older or less popular products may be less likely to be included in these materials. If so, then those products would not have an offline advertised price to enforce.

¹⁸It is worth noting that the enforcement language of Seagate’s MAP contracts changed in March 2014 to resemble more closely a minimum RPM. Nonetheless, we also compared the mean prices between MAP and non-MAP products by scraping prices in June-August 2019. When grouped by capacity level, MAP products tend to be higher priced than non-MAP products, suggesting that the difference in popularity is unlikely to be explained by a difference in prices across the 2 groups.

Table 5: Product Popularity Statistics: November 2018 Sample

Seagate Products *				
	Panelist Visits	Rating [†]	Reviews [†]	New Reviews [‡]
MAP	5.2	4.01	2,681	121
Non-MAP	1.8	3.90	1,410	86
Difference	(0.003)	(0.57)	(0.06)	(0.45)

* comScore sample contains N=75 MAP products and 21 non-MAP products in November 2018. Means and p-values calculated across all websites that consumers browsed for a given product.

[†] Scraped online August 14, 2019.

[‡] Occurring in the interval between web scrapes on June 21, 2019 and August 14, 2019.

If products selected for MAP differ in demand, the direction of bias in the $uMAP \times overlap$ coefficient depends on the average differences in the slope of the demand across the two groups. For example, if the difference in valuations l and h is smaller on non-MAP products (demand is more elastic), the $uMAP \times overlap$ coefficient would be biased downward.

6 Conclusion

In this paper, we have empirically documented the role of a widely used vertical informational restraint on the pricing of online retailers. We have shown through descriptive patterns, like the probability of retailer ties above and below the MAP level, that the recent theoretical literature which models these restraints by emphasizing consumer search appears to fit the data well. Formally, by comparing retailer pricing on products that did and did not face the restraint, a price discrimination hypothesis explained the variation in retail pricing decisions. The coefficient estimates from these models imply that the MAP policy was associated with a small but significant economic impact, and possibly enhanced retailers' ability to price discriminate.

These findings help gather early evidence about how manufacturers, retailers, and consumers interact online with vertical restraints and search costs. Future studies can build on this evidence to better understand how manufacturers use MAP to influence the profitability of themselves, their upstream rivals, and retailers. Such studies could make use of richer data (including product-level purchases) and a more flexible modeling approach (incorporating demand for different products and brands).

Finally, we note that it is important to study the effect of MAP on competition between online sellers in different types of markets, using the most recent and detailed sources of data. This is because the nature of strategic behavior that these retailers use to compete may continue to change rapidly. These changes may come from technological improvements such as pricing algorithms and individualized targeting, or from evolving legal and policy constraints.

References

- Asker, John, and Heski Bar-Isaac.** Forthcoming. “Vertical Information Restraints: Pro- and Anti-Competitive Impacts of Minimum Advertised Price Restrictions.” *The Journal of Law and Economics*.
- De los Santos, Babur, and Matthijs R. Wildenbeest.** 2017. “E-book pricing and vertical restraints.” *Quantitative Marketing and Economics*, 15(2): 85–122.
- De los Santos, Babur, In Kyung Kim, and Dmitry Lubensky.** 2018. “Do MSRPs Decrease Prices?” *International Journal of Industrial Organization*, 59: 429–457.
- European Commission.** 2017. “Final report on the E-commerce Sector Inquiry.” http://ec.europa.eu/competition/antitrust/sector_inquiry_final_report_en.pdf.
- Gilligan, Thomas W.** 1986. “The competitive effects of resale price maintenance.” *The Rand Journal of Economics*, 17(4): 544–556.
- Gorodnichenko, Yuriy, Viacheslav Sheremirov, and Oleksander Talavera.** 2018. “Price setting in online markets: does it click?” *Journal of the European Economic Association*, 16(6): 1764–1811.
- Hunold, Matthias, Reinhold Kesler, Ulrich Laitenberger, and Frank Schlutter.** 2018. “Evaluation of best price clauses in online hotel bookings.” *International Journal of Industrial Organization*, 61: 542–571.
- Ippolito, Pauline M.** 1991. “Resale Price Maintenance: Empirical Evidence from Litigation.” *The Journal of Law and Economics*, 34(2): 263–294.
- Ippolito, Pauline M., and Thomas R. Overstreet.** 1996. “Resale Price Maintenance: An Economic Assessment of the Federal Trade Commission’s Case Against the Corning Glass Works.” *The Journal of Law and Economics*, 39(1): 285–328.
- Israeli, Ayelet, Eric T. Anderson, and Anne T. Coughlan.** 2016. “Minimum Advertised Pricing: Violation in Competitive Retail Markets.” *Marketing Science*, 35(4): 539–564.

- Koulayev, Sergei.** 2011. “Search for differentiated products: identification and estimation.” *The Rand Journal of Economics*, 45(3): 553–575.
- Lafontaine, Francine, and Margaret Slade.** 2008. “Empirical Assessment of Exclusive Contracts.” In *Handbook of Antitrust Economics*. Chapter 10. Paolo Buccirossi (ed.).
- MacKay, Alexander, and David A. Smith.** 2014. “The Empirical Effects of Minimum Resale Price Maintenance.” Working Paper.
- Ornstein, Stanley I., and Dominique M. Hanssens.** 1987. “Resale Price Maintenance: Output increasing or restricting? The case of distilled spirits in the United States.” *The Journal of Industrial Economics*, 36(1): 1–18.
- Oxera Consulting and Accent.** 2016. “Vertical restraints: new evidence from a business survey.” https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/511332/Final_report_on_vertical_restraints_240316.pdf.

A Model SPE and proofs with 2 Retailers

A.1 Equilibrium without MAP

Consider the case with 2 retailers. Assume that $\delta_1 > \delta_2$, and that all parameters $\{l, h, \delta_i, \sigma_{12}, \eta, \lambda\}$ lie strictly between 0 and 1. Below, we characterize the SPE in pure strategies given these conditions, and show that it is “generally unique.”¹⁹

$$\left\{ \begin{array}{ll} \max\{\Pi^M(l, l; w^*), \Pi^M(h, l; w^*), \Pi^M(h, h; w^*)\} & ; \text{ if } l < h \leq l \frac{\delta_2 + \frac{\sigma}{2} + \frac{\eta}{2}}{\delta_2(1-\lambda^L)} \\ \max\{\Pi^M(h, l; w^*), \Pi^M(h, h; w^*)\} & ; \text{ if } l \frac{\delta_2 + \frac{\sigma}{2} + \frac{\eta}{2}}{\delta_2(1-\lambda^L)} < h \leq l \frac{\delta_2 + \frac{\sigma}{2} + \frac{\eta}{2}}{\delta_1(1-\lambda^L)} \\ \Pi^M(h, h; w^*) & ; \text{ if } h > l \frac{\delta_2 + \frac{\sigma}{2} + \frac{\eta}{2}}{\delta_1(1-\lambda^L)} \end{array} \right.$$

Depending on the value of h relative to l , under these conditions, three downstream Nash equilibria satisfy the individual rationality and incentive compatibility constraints: (l, l) , (h, l) , and (h, h) . The expressions for the manufacturer profit at the optimal per-unit rate, w^* , as well as w^* are given below.

$$\Pi^M(l, l; w^*) = l(\sigma + \eta + 2\delta_2) - (\delta_1 - \delta_2) \left(\frac{h\delta_2(1 - \lambda^L) - l(\delta_2 + \frac{\sigma}{2} + \frac{\eta}{2})}{\delta_2\lambda^L + \frac{\sigma}{2} + \frac{\eta}{2}} \right) \quad (1)$$

$$w^* = \frac{h\delta_2(1 - \lambda^L) - l(\delta_2 + \frac{\sigma}{2} + \frac{\eta}{2})}{-\delta_2\lambda^L - \frac{\sigma}{2} - \frac{\eta}{2}} \quad (2)$$

If the difference $(\delta_1 - \delta_2)$ is small, then M obtains most of the available surplus, given by the first expression. As $(\delta_1 - \delta_2)$ increases, the share of the surplus obtained by M relative to the retail surplus declines. This follows because retailer profits become more asymmetric.

¹⁹For notational ease, we depict a downstream equilibrium by the price charged by retailer 1, followed by retailer 2's price. We also omit the subscript on σ_{12} .

The optimal per-unit rate w^* decreases in $(h - l)$: retailer deviation to h becomes more attractive, requiring M to reduce w^* to induce incentive compatibility at the uniformly low equilibrium. Similarly, w^* decreases in the share of firm 2's loyalists who have high valuation, $\delta_2(1 - \lambda^L)$.

$$\Pi^M(h, h; w^*) = h [\sigma(1 - \lambda^S) + \eta(1 - \lambda^A) + (\delta_1 + \delta_2)(1 - \lambda^L)] \quad (3)$$

$$w^* = h \quad (4)$$

M sets $w^* = h$ to obtain the full surplus at a uniformly high price, when only high valuation consumers buy the product.

$$\Pi^M(h, l; w^*) = h\delta_1(1 - \lambda^L) + l(\delta_2 + \sigma + \eta) - [\max \{\Pi_1(w^*) - \Pi_2(w^*), \Pi_2(w^*) - \Pi_1(w^*)\}] \quad (5)$$

$$w^* = \frac{l(\delta_1 \frac{\sigma}{2} + \frac{\eta}{2}) - h\delta_1(1 - \lambda^L)}{\delta_1 \frac{\sigma}{2} + \frac{\eta}{2} - \delta_1(1 - \lambda^L)} \quad (6)$$

In the price dispersion equilibrium (h, l) , M obtains the total surplus when the retailer surplus (net of the per-unit rate) is equal $(\Pi_1(w^*) = \Pi_2(w^*))$.²⁰

Finally, observe that the profit expressions (1), (3), and (5) will not coincide except for precise values of the vector of willingness-to-pay parameters $(\lambda^S, \lambda^A, \lambda^L)$. This means that the SPE in pure strategies is “generally unique.”

²⁰The condition under which retailer surplus is equal is $(\delta_1 - \delta_2) = \frac{\sigma}{2} + \frac{\eta}{2}$. This is satisfied for $\delta_1 = (\delta_2 + 1)/3$, e.g. $\delta_1 = 0.4$ and $\delta_2 = 0.2$.

A.2 Proof of Proposition 1

Proof. We will consider the derivative of the manufacturer's profit with respect to σ . First, consider the 2 price uniform SPE candidates.

$$\frac{\partial \Pi^M(l, l; w^*)}{\partial \sigma} = l - (\delta_1 - \delta_2) \frac{(\delta_2 \lambda^L + \sigma/2 + \eta/2)(-l/2) - (h\delta_2(1 - \lambda^L) - l(\delta_2 + \sigma/2 + \eta/2)(1/2))}{(\delta_2 \lambda^L + \sigma/2 + \eta/2)^2} \quad (7)$$

$$= l - (\delta_1 - \delta_2) \left(\frac{\delta_2(1 - \lambda^L)(l - h)}{2(\delta_2 \lambda^L + \sigma/2 + \eta/2)^2} \right)$$

$$\frac{\partial \Pi^M(h, h; w^*)}{\partial \sigma} = h(1 - \lambda^S) \quad (8)$$

A sufficient (but not necessary) condition for [eq. \(7\)](#) > [eq. \(8\)](#) is that $h \leq \frac{l}{1 - \lambda^S}$. Next, consider the price dispersion SPE candidate. This derivative is evaluated in two different ways depending on which retailer has higher profit. Consider first $\Pi_1 > \Pi_2$.

$$\begin{aligned} \frac{\partial \Pi^M(h, l; w^*)}{\partial \sigma} &= l - \frac{\partial}{\partial \sigma} [(l - w^*(\sigma))(\delta_1(1 - \lambda^L) - \delta_2 - \sigma - \eta) + (h - l)\delta_1(1 - \lambda^L)] \\ &= l + w^*(\sigma) + \frac{\partial w^*}{\partial \sigma} (\sigma + \eta + \delta_2 - \delta_1(1 - \lambda^L)) - l \\ &= w^*(\sigma) + \frac{\partial w^*}{\partial \sigma} (\sigma + \eta + \delta_2 - \delta_1(1 - \lambda^L)) > 0 \end{aligned} \quad (9)$$

Consider next $\Pi_1 < \Pi_2$.

$$\begin{aligned} \frac{\partial \Pi^M(h, l; w^*)}{\partial \sigma} &= l - \frac{\partial}{\partial \sigma} [(h - l)\delta_1(1 - \lambda^L) - (l - w^*(\sigma))(\delta_2 + \sigma + \eta - \delta_1(1 - \lambda^L))] \\ &= -w^*(\sigma) - \frac{\partial w^*}{\partial \sigma} (\sigma + \eta + \delta_2 - \delta_1(1 - \lambda^L)) < 0 \end{aligned} \quad (10)$$

When $\Pi_1 < \Pi_2$, the marginal effect of increasing σ on $\Pi^M(h, l; w^*)$ is negative because retailer 2 captures a higher mass of price searchers. This increases the asymmetry in retailer profits, which decreases M 's fixed fee $T(w^*)$. Finally, consider that as σ increases, $\Pi_1 < \Pi_2$.

□

B Numerical Simulations of MAP Selection

In this section, we use the model to explore possible mechanisms by which the manufacturer chooses to use a MAP on a product. We do so by examining the effect of MAP on manufacturer profits in two different ways.

Figure 4: Manufacturer MAP profit change, by heterogeneity level

Figure 4 displays two panels corresponding to the difference in profit levels, with MAP relative to without. In each panel, one bar corresponds to one point on the x-y grid above. Panel (a) takes $\alpha = 1$: the “low variance” case displayed above. Panel (b) takes $\alpha = \frac{1}{3}$: a “high variance” case. The comparison between panels indicates that when variance is high, MAP is profitable at more grid points.

This figure reinforces the importance of retailer heterogeneity. In price dispersion equilibria, which characterize most of the parameter space in these simulations, retailer profits are in general unequal.²¹ By distributing advertised searchers η evenly across all retailers, the difference $(\max\{\Pi_i\} - \min\{\Pi_i\})$ tends to

²¹In contrast, when h is large relative to l , M induces retailers to price h by setting its per-unit rate $w^* = h$. This results in uniformly zero retailer profits and gives M the full surplus.

Figure 5: Manufacturer Profit Differences with MAP, High Variance Simulation

fall. Consequently, the fixed fee T^* increases, because it cannot exceed the minimum retailer profit net of per-unit rate w . On the other hand, the same search friction from MAP reduces each retailer's incentive to price l rather than h . To maintain a price dispersion solution, M must charge the marginal retailer a lower w^* . The higher T^* and lower w^* push net Π^M in opposite directions, creating variation in sign.

To complete the explanation, observe that a higher variance parameter creates more heterogeneity across retailers, which makes non-MAP profits more heterogeneous. This, in turn, raises the marginal benefit to M of a higher T^* , and renders it more profitable to trade off a lower w^* . This is the trade that MAP allows M to make.

To illustrate another possible selection mechanism, [fig. 5](#) decomposes the data plotted in the high variance case of [fig. 4](#). Each of the four panels contains a scatterplot of the mean manufacturer percentage profit gain (or loss) from imposing MAP. The panels hold constant the total share of loyal customers, $\sum_i \delta_i$, while varying the total share of price searchers, $\sum_{i,j} \sigma_{ij}$.

The smoothed fit line shows that as the total share of price searchers increases, MAP tends to become less profitable for the manufacturer. This is because at price dispersed equilibria, MAP generally provides retailers with a greater incentive to price h . Because the model takes price searchers to have l valuation,

they drop out of the market at a higher price, which tends to reduce the total surplus in the market. The figure suggests that if there is a negative correlation between search intensity and valuation, then MAP is more likely to be used on products for which consumers search fewer websites.

C Tables and Figures

Table B1: Website Visits by Session, 2011-2013 *

Website	Percent Session Visits	Total Session Visits
Amazon	73.57	1,909,085
Walmart	15.65	406,193
Best Buy	6.82	176,864
Newegg	1.97	51,075
TigerDirect	1.00	26,067
Fry's	0.47	12,296
Rakuten	0.21	5,570
Micro Center	0.21	5,399
CDW	0.07	1,788
Insight	0.03	674

* Across all 3 months in the sample

Sample comScore overlap probabilities, January 2012										
	Amazon	BestBuy	CDW	Frys	Insight	Microcenter	Newegg	Rakuten	TigerDirect	Wal-Mart
Amazon		.135	.002	.011	.	.005	.041	.001	.025	.244
BestBuy			.001	.007	.	.003	.018	.	.012	.082
CDW				.	.	.	.001	.	.001	.001
Frys					.	.	.001	.	.	.006
Insight						.	.	.	.	.
Microcenter							.002	.	.001	.003
Newegg								.001	.011	.022
Rakuten									.	.001
TigerDirect										.015
Wal-Mart										

Notes:

Each cell represents the probability that a user visited both websites in the month
From the set of all panelists, including those who did not visit an electronics-specific website
Yellow represents websites with >= 5% overlap
Grey represents websites with 1-5% overlap

Figure B1: Visitation Overlap Matrix, January 2012

Percentage of comScore panelists with Single Website Visit in Month						
	All sites [†]			Non-electronics sites [‡]		
	2011	2012	2013	2011	2012	2013
Amazon	0.470	0.523	0.574			
BestBuy	0.070	0.028	0.018	0.232	0.124	0.094
CDW	0.000	0.000	0.000	0.001	0.001	0.001
Frys	0.002	0.001	0.001	0.006	0.005	0.005
Insight	0.000	0.000	0.000	0.001	0.000	0.000
Microcenter	0.001	0.000	0.000	0.003	0.002	0.002
Newegg	0.006	0.004	0.002	0.021	0.014	0.010
Rakuten	0.000	0.000	0.001	0.000	0.000	0.005
TigerDirect	0.004	0.003	0.002	0.014	0.010	0.008
Wal-Mart	0.131	0.086	0.064			
Total	0.68	0.65	0.66	0.28	0.16	0.12

Each cell represents the probability that a user visited only that website in the month

[†]From the set of all panelists

[‡] From the set of panelists who visited at least 1 site other than Amazon, Wal-Mart in month

Figure B2: Vector of Single Site Visitors

Table B2: Multi-site Visit Probabilities
for High-Duration Panelists

All Websites		Duration (sec)		
		25 th	50 th	75 th
			60	362
		$Pr[v_{pt} \geq 2 \max\{dur_{pt}\} \geq 362]^\dagger$		
		$t = day$	$t = week$	$t = month$
μ		0.100	0.181	0.318
SE		(0.041)	(0.052)	(0.118)
N				
Electronics Websites		Duration (sec)		
		25 th	50 th	75 th
			58	351
		$Pr[v_{pt} \geq 2 \max\{dur_{pt}\} \geq 351]^\ddagger$		
μ		0.124	0.183	0.274
SE		(0.038)	(0.050)	(0.104)
N				

[†] μ , SE , N include only panelist-time observations in which at least one session on one site over the relevant interval (day, week, month) lasts as long as the 75th percentile of duration over all sessions at all sites.

[‡] Excludes Amazon and Walmart. Duration threshold is 75th percentile over all sessions at sites excluding Amazon and Walmart.