◇ *PRELIMINARY DRAFT* ◇

An Empirical Analysis of Minimum Advertised Price Restrictions*

Danial Asmat† Chenyu Yang‡

January 15, 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 dampens the effect of retailer competition on prices, whereby prices are up to 6% more dispersed than they otherwise would be. This is consistent with a model of inter-retailer price discrimination. Lastly, we also use these patterns to calibrate a simple model of consumer search with advertised prices. The model identifies under what conditions the absence of MAP would have decreased the price distribution for a given product, and by how much.

JEL Classification: L41, L81, D83

Keywords: search cost; vertical restraint; advertised price; antitrust; electronic retail

*We thank seminar participants at Michigan Ross and the 12^{th} Digital Economics Conference at Toulouse for comments, and the University of Rochester-Simon for funding. 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 Rochester, Simon School of Business: chenyu.yang@simon.rochester.edu

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 comparision 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. First, we specify an intrabrand model of how MAP impacts retailer prices when there are search costs. The model predicts that, relative to the absence of the vertical policy, MAP reduces the degree to which increased retail

[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.

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

competition cuts downstream prices. Second, using daily online retail prices of hard disk drives (HDDs), and a consumer panel measuring browsing across retail websites, we document descriptive patterns that point to the existence of search costs tied to retailer advertised price. Third, we test this hypothesis using a reduced form model that contains a set of treated (MAP) products and non-treated (non-MAP) products that are matched on observable characteristics. Finally, we use these patterns to motivate a simple calibration exercise of intrabrand pricing with a MAP. The calibration uses aggregated browsing data to proxy for consumer search parameters, and is intended to help understand how prices and profits could change in a counterfactual simulation without MAP.

We begin by specifying a version of the theoretical model of price discrimination under MAP developed by Asker and Bar-Isaac (2017) (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 habits and preferences for each retailer. We also assume a link between search type and valuation, whereby lower search cost consumers tend to have a lower valuation for the product. This structure yields pure strategy equilibria with price dispersion even in the absence of a vertical policy, which matches the pattern in the data.

We use this model to numerically simulate the conditions under which MAP is used, as well as its marginal effect on the relationship between retailer competition and product pricing. Specifically, the model implies that additional browsing overlap between sites ("competition") tends to reduce the price difference between them. We show that this effect can be *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.

Descriptive analysis reveals a number of patterns consistent with the features of the model. There are two particularly important findings that motivate a model with search frictions tied to advertised prices.

First, 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. Second, 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.

To test our primary hypothesis, we utilize a set of prices from Seagate "control" products: other 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 provides a proxy 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 leads to 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 causal 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 price effects in the online channel. A study of the total welfare implications of the policy would weigh price against service provision, on- and off-line.[3] 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. Because our price and browsing data lack associated purchase information, our model considerably simplifies the link between search type and valuation. Moreover, to focus on the main incentives of MAP, we assume a simplified unit demand function with only two discrete consumer valuations. Ongoing extensions will specify a richer demand function with additional valuations.

[3] Costly 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. Like other recent studies of vertical policies, we rely on variation in the provision of the policy in order to study its effect on prices.[4] De los Santos, Kim and Lubensky (Forthcoming) (MSRP), De los Santos and Wildenbeest (2017) (agency pricing), and Hunold et al. (Forthcoming) (MFN clauses) use variation in the *legality* of a vertical provision to identify its price effect. Our identification strategy instead exploits the presence of treated and untreated products, while attempting to control for selection effects.[5]

Our study can also be read with reference to the older literature on minimum (RPM). This is because the inter-retailer price discrimination effect of MAP contrasts with the fixed price floor of minimum RPM, distinguishing their use. We are not aware of recent empirical studies of RPM in the economics literature (Gilligan (1986); Ornstein and Hanssens (1987); Ippolito (1991); and Ippolito and Overstreet (1996)). For a discussion of the empirical challenges to studying RPM and other vertical restraints, see Lafontaine and Slade (2008).

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 the model, numerical simulations, and the main hypothesis of interest. Section 4 presents descriptive evidence that suggests the role of search costs in the market. Section 5 presents the difference-in-differences model and associated robustness checks. Finally, section 6 presents preliminary calibration results on a single product with a counterfactual simulation of no MAP policy. Section 7 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] As the cheapest major type of data storage device for computers, HDDs have long taken

[4]Theoretical models analyzing other vertical practices with search include Sandro and Sheleiga (2015), Lubensky (2017), and Garcia et al (2017).

[5]*Leegin* (2007) replaced per se illegality of minimum RPM policies with rule of reason analyis at the federal level in the U.S. It may have also strengthened the confidence associated with MAP. MacKay and Smith (2014) conduct a cross-industry difference-in-differences analysis that uses this variation.

[6]In 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%).

on a prominent role within this category. The multibillion dollar market 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.

Seagate also imposed a minimum advertised price policy on its U.S. retailers dating from at least 2009. 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.[8] 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, maximizing its service quality to its manufacturer and retail clients.

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

[7] We are not aware of a model that rationalizes the manufacturer's decision to tie MAP to MSRP, as commonly observed. This behavior may deserve further theoretical inquiry.

[8] The company was later sold to a larger entity in 2013.

directly with the manufacturer. As such, it excludes prices for products available on "marketplaces", which 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 several variants of Seagate's flagship *Barracuda* series, oriented toward workstations and high performance PCs. It also includes products from the *Momentus* series, geared toward laptops and mainstream desktop systems. Table 1 summarizes these nine products, which include frontier as well as lower storage capacity HDDs. 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
Momentus LP 250 GB	1	54.99	3
Momentus 320 GB	2	77.47	2
Momentus LP 320 GB	3	63.98	5
Momentus LP 500 GB	3	83.29	5
Momentus 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 transaction price to consumers, such as 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 leading 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 return to these data in section 4, after first presenting the formal model and numerical simulations.[9]

3 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, and there is an assumed link between those two features such that consumers who search more have a lower valuation for the product. In this type of environment, AB show that advertised price restrictions can facilitate price discrimination between retailers.

3.1 Preliminaries

3.1.1 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

[9] In our correspondence with comScore, its representatives stated that it does not retain any site-level data for its consumer sample beyond a two year window.

together, the terms of the offer (w, T) are assumed to be equal across all n retailers. The fixed fee T can be 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) $P_i^a \geq P^{MAP}$ Perfect MAP compliance

(A2) $P_i^a \geq P_i$ 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.

3.1.2 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*. Like AB, we take consumers to be non-strategic in that their search type is given exogenously. Unlike AB, we permit consumers to take on a variety of search habits 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, $Pr\left[s_i\right]$, is given below:

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

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*:

$$\begin{cases} Pr\left[\sigma_{ij} = l\right] = 1 & \forall \;\; i, j \\ Pr\left[\eta = l\right] = Pr\left[\delta_i = l\right] = \lambda & 0 \leq \lambda < 1 \end{cases}$$

To facilitate empirical calibration, this structure makes simplifying assumptions that can later be relaxed. Price searching customers σ_{ij} are assumed to be uniformly low valuation. In contrast, loyal and advertising customers η and δ_i are assumed to be low valuation with some probability less than one. This willingness-to-pay is restricted equal between both types of consumers for each i. The structure implies that in order to profitably price h, any i must attract at least some consumer types η or δ_i.

3.1.3 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).[10] 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 parameter λ. 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 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.

3.1.4 Solving for Equilibrium

M solves the following constrained optimization problem:

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

At the retailer stage of the game, there are 2^n possible combinations of downstream prices. Given a per-unit rate w, retailer price vectors must constitute a Nash equilibrium:

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

M sets the optimal fixed fee $T^*(w) = \text{argmin}\{\Pi_i \mid q_i > 0\}$. 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.

[10] If 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 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. It can be shown that the SPE in pure strategies is unique whenever searchers exhibit some asymmetry across retailers: $\delta_i \neq \delta_j$ for some $i \neq j$; or $\sigma_{ij} \neq \sigma_{jk}$ for some $i \neq j \neq k$.

3.2 Numerical Simulations

It is important to convey the important elements of the model, both to motivate the empirical analysis and to understand threats to identification. We present illustrative results from two types of numerical simulations: the effect of MAP on prices and profits, and the effect of an increase in competition on price levels.

3.2.1 Pricing and Profitability

To begin, it is useful to understand how the distribution of search types shapes equilibrium price levels, with and without MAP. To be concrete, define the set $\mathcal{S} = \{0.05, 0.10, \ldots, 0.95\}$. Then we take the x-y grid to be all pairs $\left(\sum_i \delta_i, \sum_{i,j} \sigma_{ij}\right)$ in the set $\mathcal{T} = \{x, y \ \in S \ : \ x + y \leq 1\}$. 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: $\sum_{d=1}^{D} \sum_{i=1}^{n} \frac{z_{id}}{nD}$.

Consider a Monte Carlo simulation over the retailer price vector, and take the exogenous parameters $l = 5$, $h = 6$, and the willingness-to-pay parameter $(1 - \lambda) = 1$. These parameters are chosen to illustrate a range of price equilibria across the parameter space, and thereby highlight some differences between the

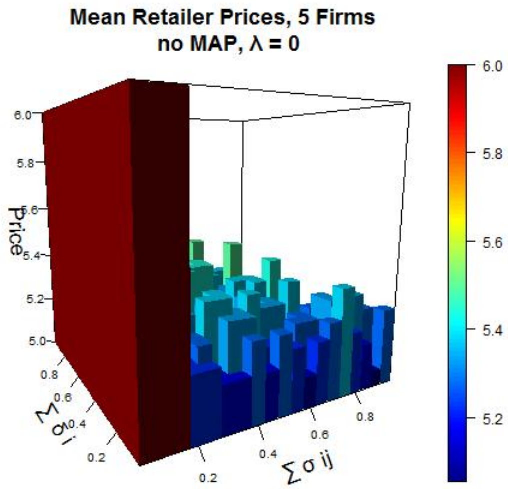

(a)

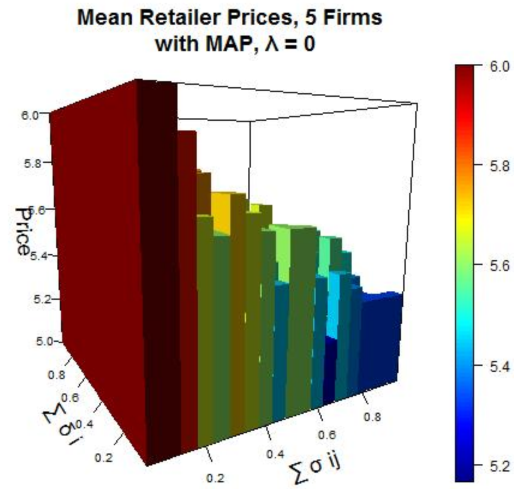

(b)

Figure 1: Mean Retailer Price Levels, $n = 5$ and $d = 20$ draws, low variance

MAP regime and default regime.[11] Figure 1 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 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.

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.

[11] As $(1 - \lambda)$ decreases, then more equilibria display uniformly low prices.

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*.[12] 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 a pervasive phenemenon in the Seagate HDD data summarized below, this effect will be important to consider in the empirical analysis.

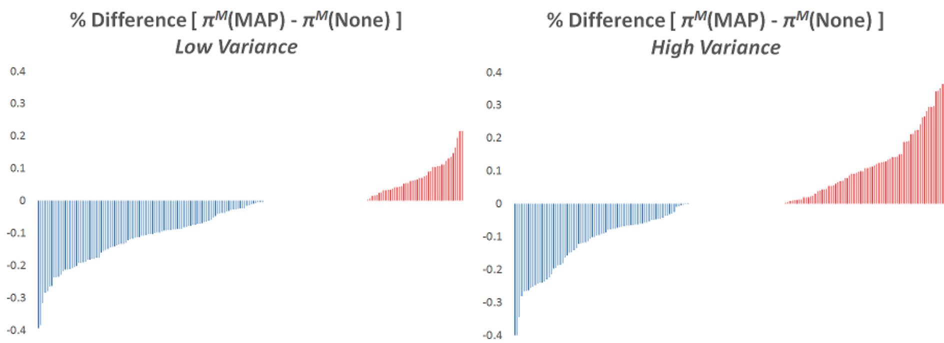

(a) (b)

Figure 2: Manufacturer MAP profit change, by heterogeneity level

It is also necessary to understand the effect of MAP on manufacturer profits. Figure 2 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.[13] By distributing advertised searchers η evenly across all retailers, the difference $(\max\{\Pi_i\} - \min\{\Pi_i\})$ tends to fall. Consequently, the fixed fee T^* increases, because it cannot exceed the minimum retailer profit net of

[12] Although 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 equlibrium.

[13] This is true in spite of M endogenously selecting w. In contrast, observe that when M wishes to induce retailers to price h, it set its per-unit rate $w = h$ in order to obtain the full surplus. This is the only w^* in a uniformly h equilibrium, and it results in uniformly zero retailer profits.

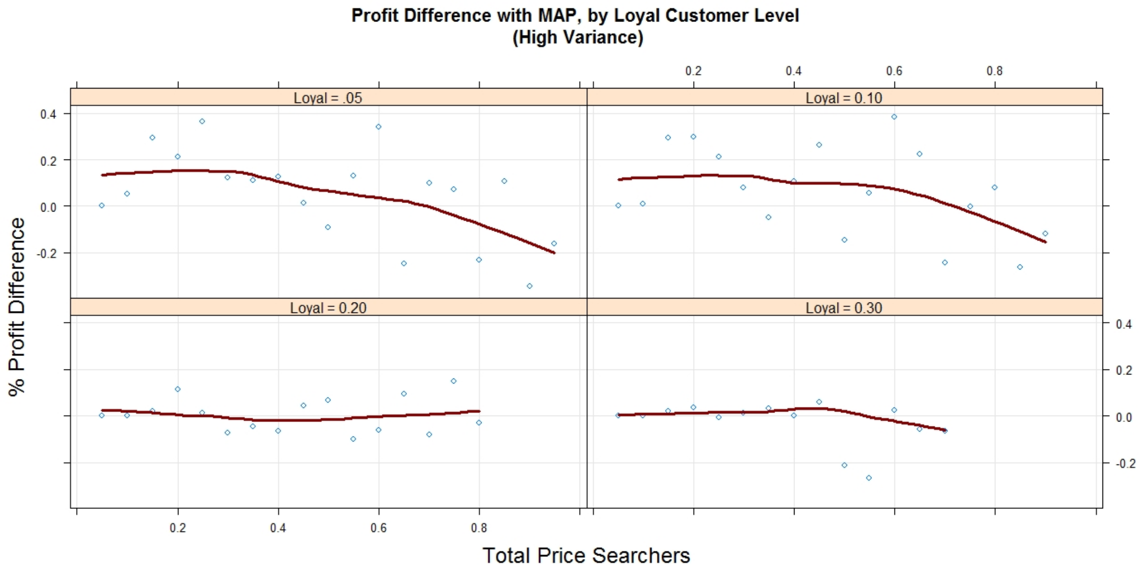

Figure 3: Manufacturer Profit Differences with MAP, High Variance Simulation

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 heterogenous. 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 comparative static, fig. 3 decomposes the data plotted in the high variance case of fig. 2. 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 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.

3.2.2 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. As before, take $l = 5$, $h = 6$, and $(1 - \lambda) = 1$. In contrast to before, take the search type probabilities $\{\eta, \delta_i, \sigma_{ij}\}$ to be deterministic rather than stochastic. Finally, let $n = 3$ firms for simplicity. The following remark will be illustrated numerically.

Remark 1. *Assume that equilibria display price dispersion, e.g. $P_i = P_j = l$; $P_k = h$.*

(i) For a positive measure of the parameter space, there exists a $\hat{\sigma}_{ij}$ such that $\forall\ \sigma_{jk} > \hat{\sigma}_{jk}$, $P_j = P_k = l$.

(ii) Moreover, $\hat{\sigma}_{jk}$ is strictly greater when the manufacturer has imposed MAP than when it has not.

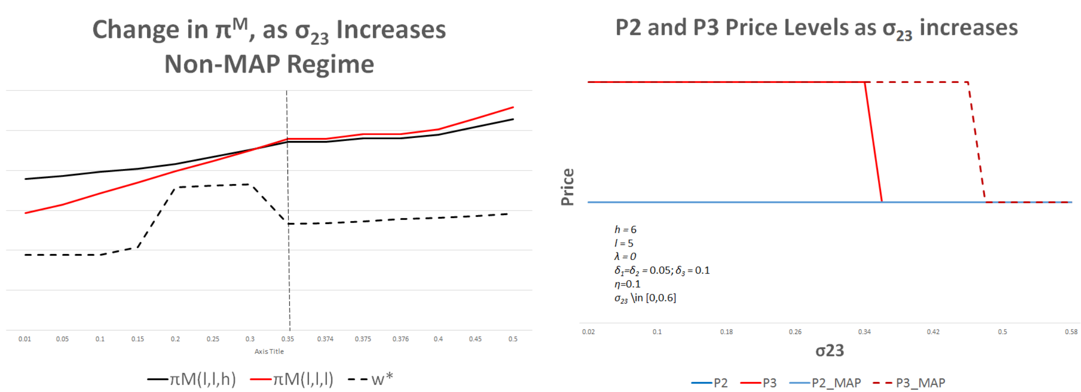

(a) (b)

Figure 4: Change in SPE as Overlap Between 2 and 3 Increases

Figure 4 graphically illustrates this remark 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_{23} 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_{23} 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, l). This equilibrium raises M's profits by equalizing the fraction of searchers σ_{23} between retailers 2 and 3, permitting a higher fixed fee T^*.

Panel (b) depicts the equilibrium prices of retailers 2 and 3 as σ_{23} increases. The blue line is constant because $P_2 = l$ in all simulations. The red solid line corresponds to P_3 in the non-MAP scenario, which switches to low price at $\sigma_{23} = 0.36$, as implied by panel (a). Finally, the red dotted line corresponds to P_3 when MAP $= h$. In contrast to the non-MAP case, the SPE does not yield a low price until $\sigma_{23} = 0.48 > 0.36$. By allocating advertised customers evenly across firms, MAP allows the manufacturer to maintain dispersed prices in the interval [0.36, 0.48].

4 Descriptive Analysis of Retailer Pricing with MAP

4.1 Product-Level Pricing

In this section, we present descriptive statistics of retail prices on MAP products, and consumer browsing habits. We use these statistics to motivate the presented model of consumer search with an advertised price restriction, including the types of consumers and the nature of retailer pricing strategies (mixed vs. pure).

To begin, we focus on a subset of the nine Seagate products subject to the MAP policy. Figure 5 depicts retailer prices for three selected products, alongside their respective MAP levels, in black. We choose these products to represent at least three different patterns that are present in the data.

First, there is dispersion across retailers and the frequency of price changes is relatively low: this is inconsistent with mixed strategy pricing.[14] Second, the MAP value is sometimes equal to the lowest retail price: this occurs frequently for the Momentus LP 320. Third, the MAP value can be matched by one retailer yet undercut by others: this occurs for several MAP levels on Momentus LP 500 and Barracuda 500.

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 MAP; or (ii) above the MAP. The last column displays the difference and t-statistic between the sample

[14] On average, 12.2 days elapse between a site changing its price. Formally, 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 25^{th} - 75^{th} percentile range of 0.01 - 0.19.

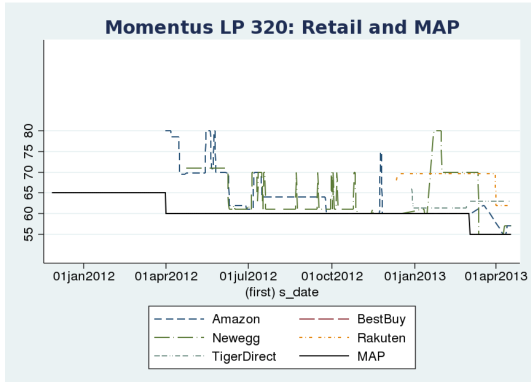

(a)

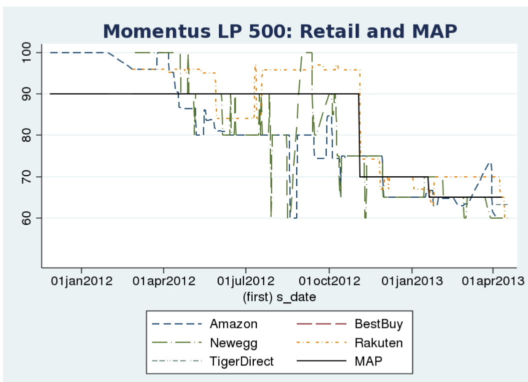

(b)

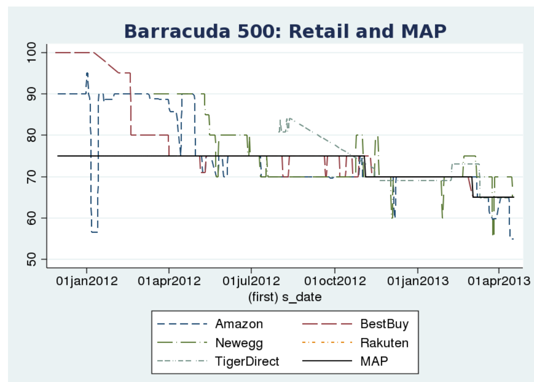

(c)

Figure 5: Retail Price and MAP Trends by Product: 2011-2013
Diagonal lines indicate stockouts

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.

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

	$P_{ij,t} <= P_t^{MAP}$ †		$P_{ij,t} > P_t^{MAP}$ †		Diff ‡	
Momentus LP 250 GB	0.92	(0.008)	0.13	(0.013)	0.79	[37.6]
Momentus 320 GB *			0.14	(0.016)		
Momentus LP 320 GB	1	(0)	0.09	(0.005)	0.91	[173]
Momentus LP 500 GB	0.31	(0.008)	0.15	(0.007)	0.16	[15]
Momentus 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; $P_t^{MAP} = \text{MAP}_t$.
‡ 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.

4.2 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.

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, $s_{pt} >= 2$, separately for three time intervals and two different website groups. The probability of visiting at least two sites is evaluated

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

	$Pr[s_{pt} \geq 2]$: All Websites †		
	$t = day$	$t = week$	$t = month$
μ	0.051	0.101	0.200
SE	(0.017)	(0.032)	(0.089)
N			
	$Pr[s_{pt} \geq 2]$: Electronics Websites ‡		
μ	0.062	0.096	0.153
SE	(0.020)	(0.044)	(0.061)
N			

† 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.

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.[15]

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 about 10%, and estimates are not significantly different across the two website groups. They rise modestly to about 20% if the panel is restricted to high-duration visits (table B2). These figures are in line with though somewhat lower than those cited by Koulayev (2014), who relates that multi-site visit probabilities in other online shopping studies with search costs have ranged from 20% - 33%. We interpret this statistic as suppport for the model's restriction that consumers visit a maximum of two websites.

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

5 MAP Effect of Competition on Retail Prices

In this section, we present empirical evidence of the effect of MAP on prices across websites. To develop this relationship, we consider a wider dataset of Seagate products: one that includes those not subject to 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.

In addition to the products subject to MAP, Seagate sold 10 others in the Momentus 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:

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

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 × 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 × 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.

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 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 - (13.5×0.17) = \$18.30; between Amazon and Wal-Mart is \$16.43.

With this in mind, the interpretaton of the interacted coefficients is straightforward. In both specifications (1) and (2), adding up the coefficients for *overlap* and $umap_1 \times overlap$ 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).

There are at least three types of endogeneity threats to interpreting this regression: non-stationarity in wholesale prices, product-level variation in consumer search probabilities, and product-level variation in demand. 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, and that the structure of the theoretical model is not invalid.[16]

[16] Clustered standard errors are relatively large given the size of the sample because prices are sticky: websites did not change prices frequently. Clustering at the website pair-product level allows arbitrary autocorrelation across days, within each website pair-product. The magnitude of the estimated variance reflects that the number of observations overstates the *de facto* size of the data generating sample.

Working from within the theoretical model, fig. 3 shows that Seagate should be more likely to use MAP on products that feature a lower mass of price searchers. If the comScore sample overlap probabilities understate the true incidence of multi-site shoppers for MAP products, then the coefficient of interest is biased upward. This possibility is challenging to address, given that the comScore data is not available below the site level in our sample period. One approach is to restrict the sample to only those sites that show the highest overlap in visitors. In estimating a sample restricted to Amazon, Best Buy, Wal-Mart, and Newegg, we estimated similar coefficients to above.

Finally, a third endogeneity possibility 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. Under this scenario, if the difference in valuations l and h is smaller on non-MAP products, the $uMAP \times overlap$ coefficient would be biased *downward*. Table B3 presents results that are consistent with this scenario, by estimating the product-time level price dispersion across MAP and non-MAP products. Using the price dispersion measures in Baye, Morgan and Scholten (2004), price dispersion is significantly higher for MAP products than non-MAP products in most capacity groups.

6 Calibration

In the final part of the paper, we conduct a calibration exercise on a Seagate product subject to the MAP policy. Because we lack disaggregated product-level browsing data, we must impose a number of additional assumptions to identify the parameters of the model. Consequently, our calibration should not be viewed as a comprehensive empirical study of MAP on Seagate's HDD prices. Rather, it is a way to use the model to further understand the effects of the policy under plausible scenarios.

6.1 Identification Assumptions

To calibrate the parameters of the model, we impose the following assumptions. First, we assume that site visitation probabilities are equal across all electronics products. This means that the comScore data in table B1 provides a measure of the matrix of price searchers on any two websites. Second, we assume that there is a product-invariant distribution of visitors between loyal and advertised searchers. This renders the vector of single-site visitors, the left panel of table B2, sufficient to allocate loyal and advertised searchers to each website. Third, we discretize retail prices into two groups, high and low. We take the most commonly selected price in each group to be the "high" and "low" valuation, respectively.

Under these (strong) assumptions, given the structure of the model, the only remaining exogenous parameter to be calibrated is the willingness-to-pay parameter λ. We observe the vector of actual retail prices from the data, $\mathbf{P}^*$.[17] We simulate the model across different λ parameters to generate predicted prices $\hat{\mathbf{P}}(\lambda)$. The identified set Λ corresponds to the matched predicted and observed prices:

$$\Lambda = \{ \ \lambda : \hat{\mathbf{P}}(\lambda) \ = \ \mathbf{P}^* \ \} \tag{1}$$

The model presented above takes all n retailers carrying each product to be exogenous. Moreover, it assumes complete information of the exogenous parameters on the part of the manufacturer. In the results below, we maintain the assumptions of complete information; however, future calibrations can accomodate weaker informational requirements.[18] They can also be modified for heterogenous search types across products, by scaling the overlap matrix of price searchers up or down.

6.2 Results

In this section, we present a calibrated result for a single Seagate product in a single year subject to the MAP policy: the Barracuda 500 GB. This product is suitable to begin calibration for several reasons. In

[17] Note that $\mathbf{P}^*$ also reveals $\mathbf{P}^{*a}$, the vector of advertised prices, given assumptions (A1) and (A2).

[18] For example, we can impose a leader-follower game such that the three most popular retailers—Amazon, Best Buy, and Wal-Mart—strategically choose prices to satisfy Nash equilibrium. In this conception, each follower would set prices with reference to the leaders, but not other followers. Because of the reduced computational burden, this conception could also accomodate more than two valuations.

2012, six retailers sold the product, which generates a suitable cross-section of firms. During that year, the MAP value stayed constant, at \$74.99. Finally, the price distribution across retailers was clustered at only a handful of values: the modal price was \$74.99, followed by \$69.99.[19] Retail prices can conveniently be rounded to either of these two values, which we take as h and l, respectively.

Table 5: Range of calibrated parameters: Barracuda 500 GB, 2012

Λ	w^*†	T^*	**Share of Man Profit**
[0, 0.05]	[65.50, 66.00]	[0.24, 0.26]	[0.92]

† Calibrated by searching intervals of 0.5

Table 5 displays the results of the calibration when the assumed distribution of loyal and advertised searchers is split evenly: $\eta = \sum_i \delta_i$. Equation 1 is satisfied when the λ parameter, which indexes the share of advertised or loyal customers with low valuation, is between 0 and 0.05. Within this range, most such customers are willing to pay the high price of \$74.99, similar to the numerical simulation. There is a corresponding range of input prices and fixed fees; the last column shows that the variation in manufacturer profit relative to total retailer profit is not large within this range.

Table 6 extends this calibration to simulate the counterfactual scenario without MAP. It displays the retail prices from the three most popular firms, along with the manufacturer profit. Interestingly, the distribution of high and low prices between the three firms remains unchanged, but Amazon switches to \$74.99 whereas Best Buy switches to \$69.99.[20] Without MAP, Best Buy is induced to reduce its price to gain more customers. This induces its rivals to raise price; Amazon raises rather than Wal-Mart because it has a larger mass of loyal customers.

This simulation shows that MAP can have an effect on manufacturer profits even without changing the distribution of prices, by changing particular retailer prices up or down. Seagate can use the different mix of customers shopping at each retailer to change its terms in order to maximize profit. In addition, the

[19]This is reflected in the Barracuda 500 graph in fig. 5. To facilitate exposition, the Barracuda 500 graph only displays prices from select retailers.

[20]The prices of remaining retailers also remain unchanged.

Table 6: Counterfactual Simulation: Barracuda 500 GB, 2012[†]

	MAP		No MAP	
	Price	**Profit***	**Price**	**Profit***
Amazon	69.99	1.44	74.99	0.96
Best Buy	74.99	0.85	69.99	0.64
Wal-Mart	69.99	0.73	69.99	0.77
Seagate	66.0[‡]	39.91	67.0[‡]	39.62

[†] Simulated price and profit parameters displayed for $\lambda = 0.05$. Simulated prices do not vary for $\lambda \in [0, 0.05]$, while simulated profits vary slightly.
[*] Profits relative to total market of unit mass.
[‡] Seagate calibrated wholesale price

simulation implies that Seagate would be worse off without MAP. This suggests that the model may be capturing some realistic elements of the market.

7 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 effectively. The coefficient estimates from these models imply that there was a small but significant economic impact of the MAP policy that weakened the effect of competition between retailers and possibly enhanced the ability to price discriminate.

These findings help gather early evidence about how manufacturers, retailers, and consumers interact online with vertical restraints and search costs. The model used to test this evidence can be extended in several directions to evaluate robustness. It can accomodate additional willingness-to-pay parameters to relax the assumed linkage between search type and valuation. By imposing further structure on downstream competition, it can also accomodate additional valuations to create a more realistic demand setting. Moroever, by conducting counterfactual experiments across different products, we can gain a better sense of the scope for different price and profitability effects of MAP in online retail.

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 within the market (technologically), or outside the market (legal and policy constraints).

References

Asker, John, and Heski Bar-Isaac. 2017. "Vertical Information Restraints: Pro- and Anti-Competitive Impacts of Minimum Advertised Price Restrictions." NBER Working Paper No. 22771.

Baye, Michael R., John Morgan, and Patrick Scholten. 2004. "Price dispersion in the small and in the large: evidence from an internet price comparison site." *The Journal of Industrial Economics*, 52(4): 463–496.

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. Forthcoming. "Do MSRPs Decrease Prices?" *International Journal of Industrial Organization*.

European Commission. 2017. "Final report on the E-commerce Sector Inquiry." http://ec.europa.eu/competition/antitrust/sector_inquiry_final_report_en.pdf.

Hunold, Matthias, Reinhold Kesler, Ulrich Laitenberger, and Frank Schlutter. Forthcoming. "Evaluation of best price clauses in online hotel bookings." *International Journal of Industrial Organization*.

Israeli, Ayelet, Eric T. Anderson, and Anne T. Coughlan. 2016. "Minimum Advertised Pricing: Violation in Competitive Retail Markets." *Marketing Science*, 35(4): 539–564.

Lafontaine, Francine, and Margaret Slade. 2008. "Empirical Assessment of Exclusive Contracts." In *Handbook of Antitrust Economics*. Chapter 10. Paolo Buccirossi (ed.).

Lubensky, Dmitry. 2017. "A model of recommended retail prices." *The Rand Journal of Economics*, 48(2): 358–386.

MacKay, Alexander, and David A. Smith. 2014. "The Empirical Effects of Minimum Resale Price Maintenance." Working Paper.

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 Appendix

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
	$Pr[v_{pt} \geq 2 \| max\{dur_{pt}\} \geq 362]$ †		
	$t = day$	$t = week$	$t = month$
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
	$Pr[v_{pt} \geq 2 \| max\{dur_{pt}\} \geq 351]$ ‡		
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.

Table B3: Price Dispersion Across MAP and non-MAP Products, Nov 2011-April 2013

	% Coef of Variation*		% Range (All)†		% Range (Min 2)‡	
Constant	0.108	(0.051)	0.148	(0.114)	0.285	(0.110)
250 GB						
320 GB	-0.045	(0.025)	-0.021	(0.062)	-0.153	(0.044)
500 GB	-0.005	(0.022)	0.047	(0.055)	-0.085	(0.037)
1 TB	-0.002	(0.011)	0.002	(0.032)	-0.013	(0.013)
1.5 TB	0.058	(0.013)	0.169	(0.035)	-0.008	(0.024)
2 TB	0.045	(0.006)	0.144	(0.018)	-0.015	(0.013)
3 TB	0.047	(0.007)	0.147	(0.020)	0.025	(0.010)
5400 RPM						
5900 RPM	0.022	(0.027)	0.014	(0.064)	0.047	(0.045)
7200 RPM	0.065	(0.023)	0.138	(0.051)	0.099	(0.033)
2.5" Form Factor						
3.5" Form Factor	-0.056	(0.025)	-0.104	(0.063)	-0.077	(0.039)
2 Retailers						
3 Retailers	-0.016	(0.013)	0.032	(0.035)	-0.044	(0.026)
4 Retailers	-0.032	(0.011)	0.017	(0.030)	-0.085	(0.019)
5 Retailers	-0.023	(0.015)	0.069	(0.038)	-0.138	(0.020)
6 Retailers	-0.026	(0.016)	0.104	(0.050)	-0.140	(0.026)
7 Retailers	-0.059	(0.024)	-0.006	(0.072)	-0.185	(0.048)
8 Retailers	-0.080	(0.023)	-0.031	(0.079)	-0.173	(0.060)
9 Retailers	-0.086	(0.021)	-0.044	(0.063)	-0.187	(0.051)
$\mathbb{1}(MAP)$	0.051	(0.008)	0.145	(0.032)	0.017	(0.025)
$\mathbb{1}(MAP)\times$ 320 GB	0.057	(0.023)	0.033	(0.058)	0.178	(0.041)
$\mathbb{1}(MAP)\times$ 500 GB	0.020	(0.023)	-0.015	(0.060)	0.096	(0.034)
$\mathbb{1}(MAP)\times$ 1 TB	-0.001	(0.012)	-0.018	(0.052)	0.044	(0.028)
$\mathbb{1}(MAP)\times$ 2 TB	0.017	(0.014)	0.065	(0.041)	0.123	(0.028)
$\mathbb{1}(MAP)\times$ 3 TB						
Month $\times$ year FE	Y		Y		Y	
N	6,899		6,899		6,899	
N Prod	19		19		19	
R^2	0.317		0.291		0.329	

* Dependent variable is σ_t/μ_t, where σ denotes standard deviation and μ mean price at t.
† Dependent variable is $(P_{it}^H - P_{it}^L)/P_{it}^L$, where superscript denotes High and Low prices for site i at time t.
‡ Dependent variable is $(P_{it}^2 - P_{it}^1)/P_{it}^1$, where superscript denotes respective ranking of second and first unique prices at t.
Standard errors clustered at product-level. Sample includes only Seagate products in Momentus and Barracuda lines that were sold by at least 2 retailers on at least 50 days. Begins from the first MAP policy effective date, November 27 2011, through April 18 2013.