Pricing with Bandits in the Long-tail: The Role of Competitive Monitoring*

Ayelet Israeli[†] Eric Anderson[‡]

July 2025

*We thank Noah Ahmadi, Liangzong Ma, and Andrew Takashi Wooders for excellent research assistance.
[†]Harvard Business School, aisraeli@hbs.edu
[‡]Kellogg School of Management, eric-anderson@northwestern.edu

Pricing with Bandits in the Long-tail: The Role of Competitive Monitoring

Abstract

Most e-commerce retailers offer a long-tail of very low demand products. Individually, these items may have low sales but collectively they are critical to the overall e-commerce business model. Because of their minimal sales, pricing is a constant challenge. The academic literature has considered price exploration as a primary source of information for price adjustments, but this approach may be insufficient in low demand situations. In this paper, we propose a bandit algorithm for long-tail products that is informed by both monitoring competitor prices and price exploration. We show that monitoring a larger competitor can inform pricing of long-tail products. Our bandit model is motivated by a unique dataset from a large e-commerce firm that regularly monitors competitor prices. We illustrate consistency between the bandit assumptions and our empirical evidence. We then show that three predictions from the bandit model are consistent with our empirical data.

Keywords: pricing, e-commerce, online retail, competitive intelligence, price monitoring, long-tail.

JEL Classification: C55, D43, M31.

1 Introduction

In e-commerce, long-tail products (e.g., Anderson 2004; 2006, Brynjolfsson et al. 2011) are both strategically and economically important. While an individual product may have relatively low sales and revenue, the collective contribution of long-tail products is essential for the e-commerce business model. It enables the concept of the "endless aisle" allowing online retailers to offer millions of products to customers, including niche or rarely purchased items. Retailers often achieve this without the need to own or stock these items, instead relying on second- or third- party logistics companies that drop ship products on demand.

A key challenge is pricing long-tail products. When demand is uncertain, the standard approach to learning about demand is price exploration, and algorithmic pricing is widely used by e-commerce retailers to vary prices and learn about demand (Aparicio and Misra 2023). Much of the academic literature has focused on learning among high-demand and popular products (e.g., Chen et al. 2016, Hansen et al. 2021, Brown and MacKay 2023, Aparicio et al. 2021; 2024), but the long-tail has received far less attention. In the long-tail, demand signals are more rare and expensive to obtain. For example, it may take many time periods for a demand signal (i.e., a purchase) to arrive, making price experiments low powered and uninformative.

In this paper, we introduce a new approach for pricing in the long-tail that combines price exploration and competitive monitoring. We propose a bandit model for the long-tail that uses these two information sources as demand signals. We also present descriptive evidence from a large e-commerce retailer that sheds light on the process of adjusting prices among long-tail products and the association with future sales. A novel contribution of our research is to demonstrate that competitive monitoring is an important source of demand information among long-tail products. While on average monitoring is informative, we show, both in our model and in our empirical analysis, that monitoring a larger competitor's prices can be a particularly important source of information for a smaller competitor. When larger competitors are trying to learn about demand through their own price exploration their prices may provide a valuable, cost-effective demand signal for a smaller retailer.

In our model, we adjust the bandit framework to lead to price changes even without observing demand. This is crucial in the long tail, where demand is rare. The model relies on key components of our institutional settings: (a) monitoring and exploration are costly activities; and (b) larger firms have a larger customer base, leading to greater demand. Reflecting this setting, our model includes a large retailer that is likely to have more frequent demand than the smaller retailer.

We note that assumption (a), costly monitoring, differs from other papers, such as Miklós-Thal and Tucker (2019), who assume that firms may have full information about competitor prices due to costless monitoring. As we will show, our assumption is motivated by extensive empirical evidence from a durable goods retailer. We find that the retailer monitors intermittently, which is consistent with costly monitoring.

The bandit model results show that the higher the demand at the large retailer (relative to the smaller one), the more effective the monitoring is for the smaller retailer. We also show that because monitoring is costly, it is not optimal for a firm to always monitor. Finally, because exploration is also costly, there is a trade-off between the signals from monitoring and exploration — when exploration becomes more expensive (while the cost of monitoring is held constant), it is optimal for the retailer to engage in more monitoring.

We utilize a unique empirical dataset from an e-commerce retailer that includes prices and sales for over one-hundred thousand products. The vast majority of these products are long-tail products that have very low demand. Importantly, we also observe the daily monitoring activity of the focal firm. Thus, we observe what the focal retailer knows about competitor prices. This contrasts with other approaches in the literature that utilize scraped, historical data and assume this is what a firm observes. We use this data for two purposes.

First, we use the empirical data to develop the assumptions for our bandit model. For example, one assumption is that monitoring is costly and in-turn this implies that monitoring may not occur every period. Another assumption (or premise) of our bandit model is that price monitoring is associated with future price adjustments, which we demonstrate with our empirical data.

Second, we generate predictions from the bandit model and investigate these in our empirical context using a quasi-experimental design. There are three main findings from the analyses. First, we show that price monitoring and price exploration are substitutes. Second, we show that monitoring the largest competitor is more beneficial to the focal firm than monitoring smaller competitors. Third, we show that increased monitoring frequency is associated with increased revenue. Overall, the predictions from the bandit model are consistent with our empirical results.

The paper continues with a review of the literature in Section 2. In Section 3, we present the bandit model and the simulation results. In Section 4, we provide institutional details about the focal firm as well as descriptive analysis of our empirical data. This section demonstrates how the data supports the assumptions for our bandit model. In Section 5 we use the data to investigate key predictions from the bandit model. The paper concludes with a discussion in Section 6.

2 Literature Review

2.1 Long-Tail Products

Our paper contributes to the literature on the long-tail, a term that was coined by Anderson (2004) referring to the idea of "selling less of more" in e-commerce, and has come to represent those items with low demand that are sold online. A number of empirical and theoretical papers have sought to understand factors that contribute to the existence of long-tail in online retail settings. Empirical papers include Brynjolfsson et al. (2011), who compare an online versus catalog setting and argue that search technology contributes to concentration. Hinz et al. (2011) study the video-on-demand market and identify various drivers of demand in the long-tail, such as seasonality and acquisition of new customers. Choi and Bell (2011) study the consumer diaper market and show that preference minorities contribute to the long-tail. Theoretically, Yang (2013) shows how firms target consumers affects the long-tail. Similarly, Quan and Williams (2018) show that the long-tail may be a consequence of demand aggregation over geographies. Finally, recommendation systems may also contribute to the existence of the long-tail (Oestreicher-Singer and Sundararajan 2012, Fleder and Hosanagar 2009).

There is no standard definition of the "long-tail" term, so we choose to define popular items as those that generate the top 80% of revenue and the long-tail as those in the bottom 20% of revenue (Clayton 2019).[1] Bar-Isaac et al. (2012) show that in some markets (e.g., books), there may be an increase in both super-star products and long-tail products. Our paper finds a similar result as roughly 1% of products represent 80% of revenue.

2.2 Price Adjustments

In the context of modern, e-commerce retailers, price adjustments are determined by algorithms. As noted by Hansen et al. (2021), "Algorithms set real-time prices for an array of products for which the retailer has incomplete demand information." Aparicio and Misra (2023) define algorithmic pricing as "automation" of retail price adjustments. They note that there may be a broad range of automated pricing algorithms, ranging from sophisticated AI algorithms to simple pricing rules like "match a competitor." In our setting, the retailer sells millions of products and therefore manual pricing is simply infeasible. We were not provided with additional insights regarding the pricing algorithms, but we were informed that prices were not set manually.

Past price exploration and historical demand are two common sources of data that may inform a retailer's price algorithm and lead to price adjustments. Many retailers now have sophisticated data science teams that deploy price experiments. For e-commerce retailers, online experiments and price exploration algorithms are extremely common. A recent paper by Cooprider and Nassiri (2023) describes how Amazon uses different types of experiments to assess demand and in-turn adjust prices.

Extant research has examined how price experiments and pricing algorithms, like multi-

[1] See `https://www.pocketbook.co.uk/blog/2019/09/03/long-tail-infinite-consumer-choice/`, accessed May 26, 2023

armed bandits, affect equilibrium prices among products with high demand (e.g., Ganti et al. 2018, Calvano et al. 2020, Hansen et al. 2021, Weaver et al. 2024). These algorithms rely on observing demand to update information about prices. Therefore, they cannot be applied to the long-tail where demand is scarce. Hansen et al. (2021) consider a market where a firm does not observe it's competitor's prices, which makes this an omitted variable in demand estimation. They show that when the signal to noise ratio of a price experiment is high then monopolistic prices may arise. In contrast, when the signal to noise ratio is low then competitive prices emerge. In our paper, long-tail products have a low signal to noise ratio, but unlike Hansen et al. (2021) the focal retailer in our study observes competitors' prices via periodic price monitoring. Thus, it is not clear that their model provides a clear prediction on equilibrium prices in settings with competitive monitoring.

A core challenge with experimentation and demand estimation among long-tail products is the scarcity of data: selling zero units is the norm. Mussi et al. (2022) note that in practical e-commerce settings it is rare to observe at least one transaction per product. The disparity among popular and long-tail products is addressed by Huang et al. (2022), who show that price learning is faster for products with more precise sales information. In related research, Adam et al. (2024) show that standard BLP-style estimators are biased when there is a long-tail of products with very low sales.

When historical demand is sparse or missing, researchers have explored obtaining new data and deploying new algorithms. For example, in the context of new products, for which sales data is missing, Cao and Zhang (2021) develop an experimental methodology that generates new data to forecast demand. To address bias in estimation, Adam et al. (2024) propose a new algorithm, which is a two step estimation strategy that relies on estimation of a neural network (step 1) and then an adjusted BLP-style demand model (step 2).

An additional source of data that may inform current price adjustments is past monitoring of competitor prices. There are a myriad of vendors who sell price tracking services (e.g., Wiser Solutions, Octoparse, Repricier.com, Intelligence Node, and Prisync are a handful of third party solutions). In addition, researchers have scraped prices of competing firms

(Cavallo and Rigobon 2016) and incorporated this in their analyses. A limitation of both scraping and third party data is that a researcher is unaware of the information set of a retailer. Absent this knowledge, one approach has been to assume that retailers have full information about each others prices (e.g., Miklós-Thal and Tucker 2019, Brown and MacKay 2023). We address this challenge in our research by obtaining data on competitive price monitoring from a focal retailer. As we will document, the focal retailer typically has limited competitive intelligence from a small number of key competitors.

Within the monitoring literature, our paper is related to Fisher et al. (2018), who conduct a field experiment with a retailer to determine the best-response price for a retailer that monitors competitors' prices, and Brown and MacKay (2023), who use high-frequency monitoring data and pricing technology to examine best-response in competitive equilibrium. Fisher et al. (2018) note that determining best-price response hinges on several factors including margins, the competing retailer characteristics, whether consumers shop around (i.e., search), and price elasticity. A key feature of their model is that a retailer may adjust price response based on these factors. They construct a field experiment to measure demand and then develop an algorithm for best-price response and document a substantial increase in profit. Both Fisher et al. (2018) and Brown and MacKay (2023) are largely focused on popular products where a retailer can obtain an accurate estimate of own-price elasticity and cross-price elasticity via price exploration. In contrast, our research focuses on the long-tail, where demand estimation is extremely noisy and difficult.

A final source of information for price adjustments is historical sales of similar products. For example, a retailer may carry a large set of products offered by a single manufacturer, like Black+Decker or Puma. Some of these products may be popular while others may be in the long-tail. Mussi et al. (2022) examine a situation analogous to this and develop an alternative approach to pricing long-tail products. The authors identify similar products via a search algorithm, aggregate sales for a group of products, and then estimate demand via a bandit model. They demonstrate the efficacy of their method through simulation and then, impressively, they show that the algorithm leads to a 90% increase in revenue for long-tail products at an e-commerce retailer.

Our research complements these papers as we demonstrate how an e-commerce retailer obtains information about long-tail demand from all of the information sources described above. Our main contribution is that we propose an adjustment of the bandit framework to enable learning about prices for low-demand products, utilizing both price exploration and competitive price monitoring. We demonstrate that price monitoring of a large competitor provides valuable information for adjusting prices of long-tail products. In turn, we show that these price adjustments are associated with more sales and revenue.

3 Bandit Model

This section presents the bandit model for long-tail products that draws on the institutional details of our empirical context, and simulation results illustrating the model's results.

3.1 Model

Consider a market with two firms, F_L and F_S, where F_L is assumed to be much larger than the small firm, F_S. Both firms sell a single product that has relatively low sales in every period (i.e., a long-tail product). Across many periods the modal sales each period is zero. Both firms sell to similar customers who all have identical valuation V for the product, but V is not known to either firm. In any period, at most one customer arrives at either firm and purchases if the price, p, is less than or equal to their valuation. At the large firm, the probability that a customer arrives and considers buying is α_L and at the small firm it is α_S, where $\alpha_L > \alpha_S$.

In period 1, both firms believe that valuation, V, is distributed uniformly between $[\mu - x, \mu + x]$ and both firms face marginal cost, c. Given this setup, the optimal price is $p^* = (\mu + c + x)/2$ if the firm prices in a single period.[2] Initially, we assume that both firms set prices at p^* if $p^* > \mu - x$, and otherwise both firms set $p^* = \mu - x$ (i.e., the lower bound of

[2]See Online Appendix A for derivation.

V) in period 1. Firms then update p^* based on the results of price exploration. To illustrate the model in our simulations, we assume that the initial value of $p^* = \mu - x$.

There is a marginal cost per period of every price experiment, $c(\beta)$, which leads to price stickiness. In each period, we assume that each firm places a mass point, $\beta_L(t)$ or $\beta_S(t)$, on a focal price $p^*(t)$ and then explores over the remaining prices. Thus, the large firm experiments with probability $1 - \beta_L(t)$ and the small firm experiments with probability $1 - \beta_S(t)$. The mass point is rationalized by the fact that exploring prices is costly. This mimics our data where firms typically charge a regular price but occasionally explore prices.

Large Firm Behavior Since the large firm is uncertain about demand, it explores prices using an algorithm, which places weight $G(\cdot)$ on each price in the support $[\mu - x, \mu + x]$. $G(\cdot)$ is a mixture of a mass point at $p^*_L(t)$ and a uniform distribution. There is a mass point $\beta_L = \beta_L(1)$ on $p^*_L(t = 1)$ in period 1. (i.e., the focal price). The large firm uniformly explores non-focal prices with probability $1 - \beta_L(1)$ in period 1. In each period t, the large firm takes a draw from $G(\cdot)$ and posts this price. If there is zero demand at this price no action is taken. In the next period the large firm takes another draw from $G(\cdot)$ and the process repeats.

If demand is realized in period t at $p_L(t)$, the large firm learns a lower bound on V. There are then two potential changes to $G(\cdot)$. First, $G(\cdot)$ is updated such that no prices less than $p_L(t)$ are searched; the support of $G(\cdot)$ is now $[p_L(t), \mu + x]$.

Second, the large firm determines whether $p^*_L(t)$ is still optimal. It is straightforward to show that $p^*_L(t + 1) = \max\{p^*_L(t), p_L(t)\}$. Thus, if demand is realized at a price that is greater than the current $p^*_L(t)$, the mass point is moved to that price.

Regardless of whether the mass point changes, the large firm increases the mass point by placing more weight on p^*. Since prices in the range $[\mu - x, p_L(t)]$ are no longer explored, the mass on $p^*_L(t)$ increases by $[p_L(t) - \min\{support(V)\}]/2x$.

Over time, the large firm learns more about possible values of V, $p^*_L(t)$ is weakly increasing, and price exploration $1 - \beta_L(t)$ is weakly decreasing. Note that the benefits to price exploration in the long-tail are purely about margin, not volume. If the large firm sets $p = \mu - x$ then a purchase is realized with certainty whenever a customer arrives. Learning about the true valuation enables the firm to charge a higher price.

Small Firm Behavior The small firm is also uncertain about demand, V, and utilizes a similar algorithm to explore prices. However, the small firm has the option to monitor the large firm's prices. Before the game starts, the small firm commits to a monitoring policy, M: (i) monitor every period (ii) monitor with probability $0 < \gamma < 1$ every period or (iii) never monitor. The marginal cost of monitoring in each period is given by $c(\gamma)$.

Within each period, we assume the following sequence of actions and outcomes for the small and large firm:

1. Both firms set prices (experimental or focal prices, p_j^*) at the start of the period.
2. The small firm monitors the large firm.
   (a) If $M = 1$ the large firm price is monitored.
   (b) If $M = \gamma$ then with probability γ the large firm price is monitored.
   (c) If $M = 0$ then no monitoring occurs.
3. Demand, $q_j(t) \in \{0, 1\}$ is realized at both firms.
4. Update $F(\cdot)$ and $G(\cdot)$.

We further assume that the small firm can rationally anticipate the pricing algorithm of the large firm. In other words, the small firm knows that the large firm is exploring prices over a similar range of prices and is trying to learn demand.

If the small firm chooses $M = 0$ (no monitoring) it utilizes a pricing algorithm that is similar to the large firm. The small firm chooses a focal price, $p_S^*(t)$, based on prior beliefs about V. We assume that while V is unknown it is within the range of searched prices $[\mu - x, \mu + x]$. The small retailer places a distribution over this range of prices, $F(\cdot)$, that has a mass point at p^* and is uniform over the range of possible valuations. In period 1, the mass point on $p_S^*(1)$ is $\beta_S = \beta_S(1)$. If $M = 0$, price exploration and updating to $F(\cdot)$ are analogous to the process for the large firm.

Now consider the case where the small firm monitors the large firm, $M > 0$, and there is no demand at the small firm. The small firm is aware that the large firm has a regular price $p_L^*(t)$, which is not known. And, the large firm also explores over a range of prices. Let $p_L(t)$ equal the current monitored price and $\mathbf{p_L}(t')$ equal the vector of all past monitored prices, where $t > t'$. If $p_L(t)$ is equal to any previous price, then the small firm learns that this is the current mass point (i.e., focal price or regular price of the large firm).

Recall that initially, both firms set the same focal price: $p_S^*(1) = p_L^*(1)$. And, the large firm may increase the focal price if demand is realized at a price greater than $p_L^*(1)$. Thus, if $p_L(t)$ is greater than the small firm's focal price, the small firm updates its focal price to $p_L(t)$. $F(\cdot)$ is updated so that the range of prices explored is now in the range $[p_L(t), \mu + x]$. Finally, the mass on is $p_L(t)$ updated.

In contrast, if the small firm learns that the large firm has a mass point on $p_L(t)$ and this price is less than the current focal price of the small firm, then there is no update to $F(\cdot)$. In this situation, the small firm has better information than the large firm about the possible values of V.

Thus, competitive monitoring provides two potential benefits to the small firm. First, unprofitable prices are not explored. Second, when demand occurs it is more likely to occur at a higher price, which increases gross margin.

Stopping Rules As firms explore prices, the marginal benefit of experimentation converges to zero as firms learn the true V. Since the marginal cost of experimentation is assumed to be non-zero, both firms rationally terminate experimentation in finite time.

We incorporate two stopping rules. Both firms stop exploration when either (1) or (2) is satisfied.

1. Cost vs value: When the cost of exploration exceeds the realized marginal value of exploration.

2. Limit case: After S periods with no update to the focal price.

When (2) is satisfied, the small firm also terminates monitoring if $M > 0$, as it exhausted its learning attempts. Details on the implementation of these stopping conditions are described in Web Appendix A.

3.2 Results

3.2.1 Illustrative Example

To illustrate how price adjustment works, we set the parameters and present the dynamic behavior of each firm in Figure 1. There are nine panels in Figure 1. The top three panels illustrate the relative speed of updating the minimum support on V for the large and small firm. The upper left panel considers no monitoring ($\gamma = 0$). The solid black line is the lower bound of the support of the large firm, which is by assumption weakly increasing. The dashed blue line shows that the small firm updates its belief about the lower bound of V more slowly than the large firm. In this example, the small firm believes the lower bound of V is less than 1.0 while the large firm correctly believes it is closer to 1.1 (the true value of V, which is illustrated by the red dotted line).

But, what happens when the small firm now monitors? Moving to the upper middle panel, we increase monitoring to ($\gamma = 0.15$), which is typical of the monitoring rates we observe in our data. Now, the small firm updates its belief about V at a faster rate and eventually price comes close to the the true value of V by the final period. Finally, in the upper right panel we consider the case there the the small firm monitors every period ($\gamma = 1$). In this scenario, the price path of both firms is nearly identical. The only gap is driven by the time it takes the small firm to differentiate a regular price from an experimental price. If monitoring costs are zero, then clearly $\gamma = 1$ is optimal.

The middle row of Figure 1 illustrates the prices of the large firm under the same three monitoring scenarios. Each open circle represents the price charged in a period by the large firm when no demand is realized. A solid dot represents a case where demand is realized but there is no update to p^* or G(). Finally, a solid red dot indicates a case where demand is realized and p^* and G() are updated. Visually, a sequence of horizontal dots that begins

Figure 1: Illustrative Example - Firm Dynamics

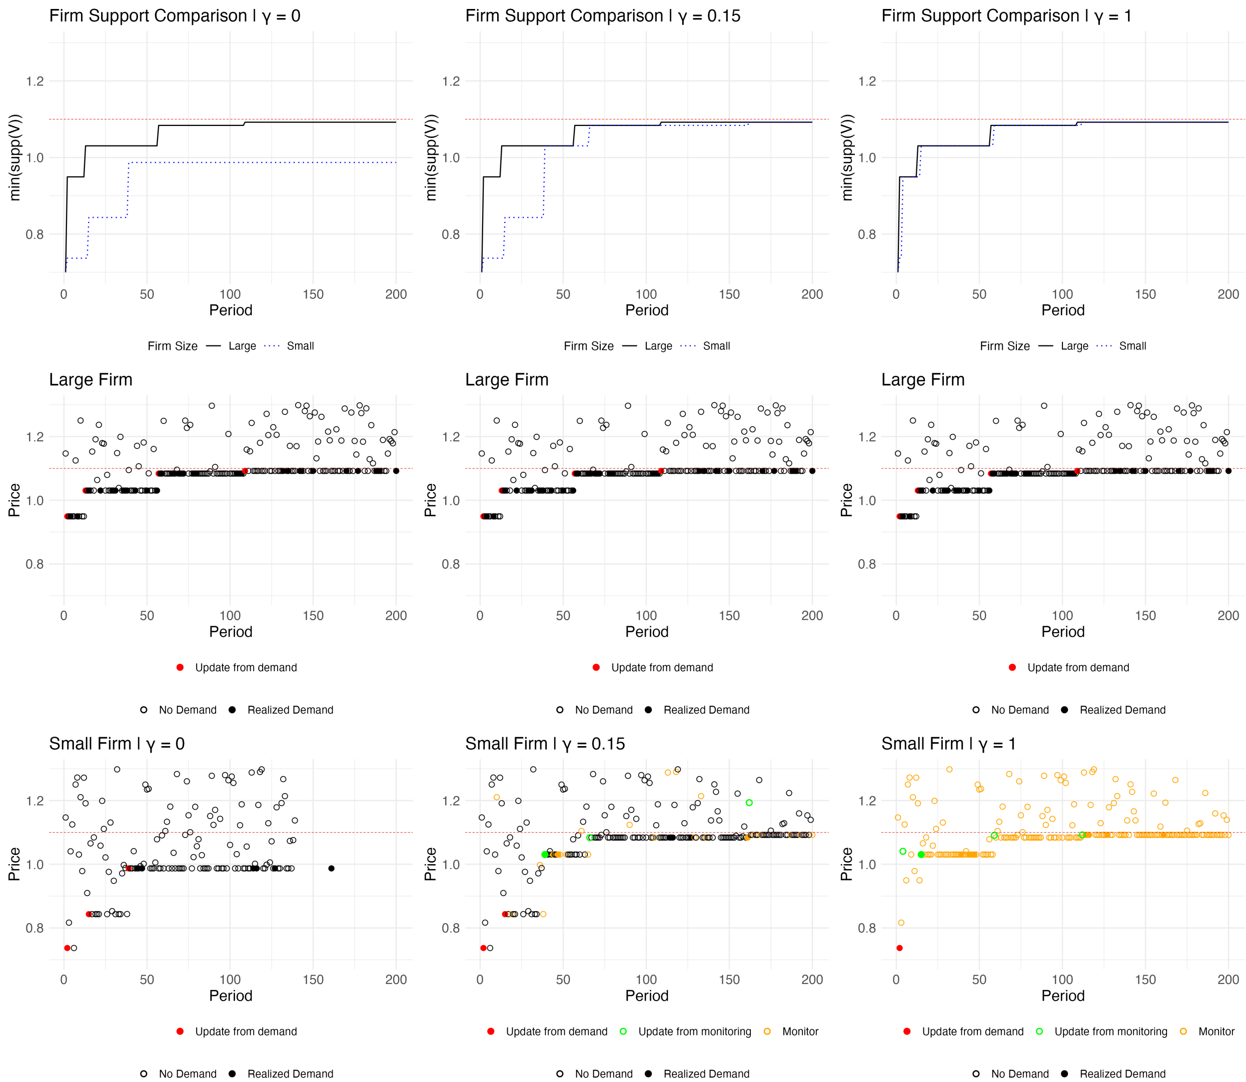

Notes: The parameters used to create this example are: $V = 1.1$, $\alpha_L = 0.25$, $\alpha_S = 0.05$, $\beta_L(1) = \beta_S(1) = 0.1$, $c = 0.05$, $x = 0.3$, $\mu = 1$,$c(\beta)_L = 0.002$, $c(\beta)_S = 0.0025$, $c(\gamma) = 0.01$, $S = 200$.

with a solid red dot reflects a regular price, p^*. Not surprisingly, the price path of the large firm is identical in all three panels since it is not affected by small firm monitoring. By roughly the 100th period of the game, the large firm has come very close to learning V via price exploration.

The bottom three panels of Figure 1 show the price dynamics for the small firm under the three monitoring scenarios. In the lower left panel, the small firm only learns via price exploration. Since demand signals are rare, the small firm never reaches the true value of V, which is 1.1. In the lower middle panel, monitoring is increased to $\gamma = 0.15$. The open yellow dots represent periods where the small firm monitored the large firm, but there was no update in beliefs about V. The open green dots represent cases of both monitoring and updating beliefs. What we observe is that small firm updates its regular price more quickly and approximates the true value of V by period 60. Finally, in the lower right panel, the small firm monitors in every period and the price path mimics the large firm.

3.2.2 Simulation Results

In Figure 1, we illustrated the features of the model for one simulation over 200 periods. For our main results, we run 1,000 simulations of the model over 500 periods while fixing the baseline parameters at the same values. The results are summarized in Table 1.

Table 1: Bandit Model Simulations

	Parametrization: $V = 1.1$, $\alpha_L = 0.25$, $\alpha_S = 0.05$, $\beta_L(1) = \beta_S(1) = 0.1$, $c = 0.05$, $x = 0.3$, $\mu = 1$, $c(\beta)_L = 0.002$, $c(\beta)_S = 0.0025$, $c(\gamma) = 0.01$, $S = 200$.							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Exploration Periods	Demand Realized	Monitoring Periods	Update to Competitor	Focal Price at Stop	Total Revenue	Total Profit	Periods taken to Stop
Large Firm	15.016%	22.020%	0.000%	0.000%	1.085	118.698	113.043	161.352
Small Firm ($\gamma = 0$)	20.761%	4.396%	0.000%	0.000%	1.016	22.134	20.775	174.438
Small Firm ($\gamma = 0.15$)	17.115%	4.364%	5.575%	0.355%	1.087	23.488	21.904	183.490
Small Firm ($\gamma = 0.50$)	15.212%	4.388%	17.678%	0.509%	1.088	23.698	21.527	174.944
Small Firm ($\gamma = 1$)	14.656%	4.398%	34.972%	0.588%	1.088	23.766	20.735	172.938

Note: This Table summarizes pricing behavior across 1,000 simulations of 500 periods each. Columns (1)–(4) are the average proportion of periods in the first 500 periods of the simulation. Column (5) is the average price. Columns (6)–(7) are cumulative of the first 500 periods. Stop indicates the period when the firms stop experimentation.

In each row of Table 1, the large firm behavior is held constant (presented in the first row), and the small firm varies it's monitoring rate. To begin our discussion, consider the scenario where there is no monitoring (Row (2)). Here, the small firm spends more time exploring than the large firm (20.8% vs 15% of the time, and Column (1)). On average, demand is realized 22% of the time for the large firm and only 4.4% of the time for the small firm (Column (2)). Because there is no monitoring, there are no monitoring periods and there is no update to competitor prices (Columns (3) and (4)). The focal price at the end of 500 periods for the large firm is on average 1.085, which is close to the true valuation 1.1. In contrast, the small firm's final focal price is on average 1.016 (Column (5)). Columns (6) and (7) report the revenues and profits, respectively. These will serve as a baseline for the monitoring scenarios we consider next. Finally, the small firm took 8% more periods than the large firm to stop exploring prices (Column (8)).

In rows 3, 4 and 5 we gradually increase the monitoring rate of the small firm from 0.15 to 1.0. In Column (1) we see that price exploration decreases as monitoring increases, which illustrates that price exploration and monitoring are substitutes. However, demand realization is not necessarily higher when there is more monitoring (Column (2)). We observe that demand realized is roughly 4.4% across the four monitoring scenarios. Again, this reinforces the intuition that in this simulation the goal of monitoring is not to generate more demand but to generate more revenue for each unit sold.

Column (3) shows how the monitoring parameter, γ, is related to the average percentage of of periods monitored. Note that even though $\gamma = 1$ leads to monitoring each period, there is a point where monitoring ends because of a stopping rule. Thus, under all monitoring scenarios the simulation assumes that monitoring eventually becomes zero.

Column (4) show that as monitoring increases, the small firm is more likely to update prices to the monitored competitor price. Because we are studying long-tail products, updates are rare events and occur less than 1% of the time in this simulation.

Column (5) shows that even with relatively infrequent monitoring (15%), the average focal price for the small firm at the end of the simulation is similar to the average stopping focal price at the large firm. Thus, even a small amount of competitive monitoring can close the gap in prices between the firms.

Column (6) shows that small firm revenue is monotonically increasing in the monitoring rate. Since units sold is not increasing with monitoring, this implies that more monitoring leads to an increase in average price paid. While small firm revenue increases with more monitoring, total profit for the small firm is concave in monitoring costs (Column (7)). Among these four monitoring scenarios, $\gamma = 0.15$ has the highest profit.

Finally, in Column (8), the average periods to stop experimenting are less than 200. One of our stopping rules was based on no information revealed over 200 periods, which rarely holds. Thus, the reason both firms stop experimenting is due to the "cost vs value" stopping rule.

3.2.3 Comparative Statics

In this subsection, we explore three variations on the results: the impact of the cost of monitoring on the optimal monitoring rate, the impact of the cost of exploration on the optimal monitoring rate, and the impact of the arrival ratio between the large and small firms on profitability.

Impact of Monitoring Cost Figure 2 presents the optimal monitoring rate for different values of the cost of monitoring. Except for varying the cost of monitoring, the monitoring rate, and the valuation V, the parameterization is the same as in Table 1. To calculate the optimal monitoring rate, we ran simulations of the model for each monitoring cost level $c(\gamma)$ and value V (601 total values for each cost level), and compute the average profits for each monitoring rate at that cost. Then, for each cost, the optimal monitoring rate is the one associated with the maximum average profits. The Figure illustrates that the higher the cost of monitoring, the lower the optimal monitoring rate. For the parameters in Table 1, the optimal monitoring rate is 0.15, which is presented in Panel B in the table.

Impact of Exploration Cost Figure 3 presents the optimal monitoring rate for different values of the cost of exploration. Except for varying the cost of exploration, the monitoring rate, and the valuation V, the parameterization is the same as in Table 1. Similarly to Figure 2, to calculate the optimal monitoring rate, we run 601 simulations of the model for each exploration cost level $c(\beta)_S$, and compute the average profits for each monitoring rate at that cost. Then, for each cost, the optimal monitoring rate is the one associated with the

Figure 2: Simulation Results - Cost of Monitoring

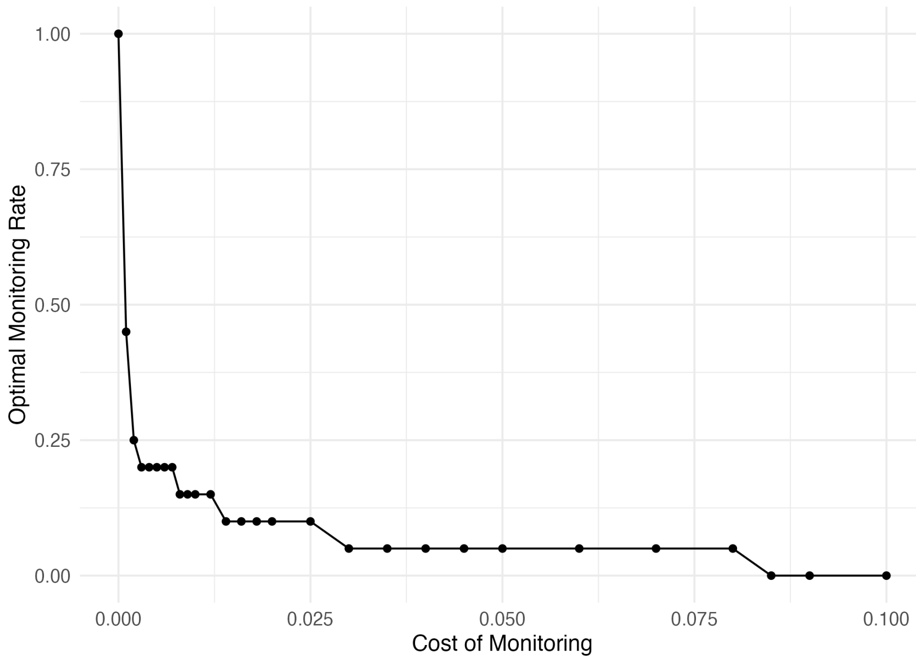

maximum average profits. The Figure illustrates that monitoring serves as a substitute to exploration. While the optimal monitoring rate is 0.15 up to a cost of exploration of 0.0055, once the exploration cost is higher, the optimal monitoring rate increases, to offset the cost of exploration.

Arrival Ratio Figure 4 presents the mean profit for three different arrival ratios. When the arrival ratio equals 1, the small and large firm have the same customer arrival rates, and when it is greater than 1, the large firm arrival rate is larger than the small firm by that factor. Except for varying the large firm arrival rates and the monitoring rates, the parameterization is the same as in Table 1. Each of the curves illustrates what we saw in the previous section — until about a monitoring rate of 0.15, more monitoring yields higher profit. We also observe that for the three arrival rates, the optimal monitoring rate is roughly the same ($\sim$0.15). Monitoring at a rate greater than 0.15 reduces profit due to the cost of monitoring. Comparing the curves to each other, as the arrival ratio increases we observe that there is greater profit for the small firm. Intuitively, monitoring a larger competitor

Figure 3: Simulation Results - Cost of Exploration

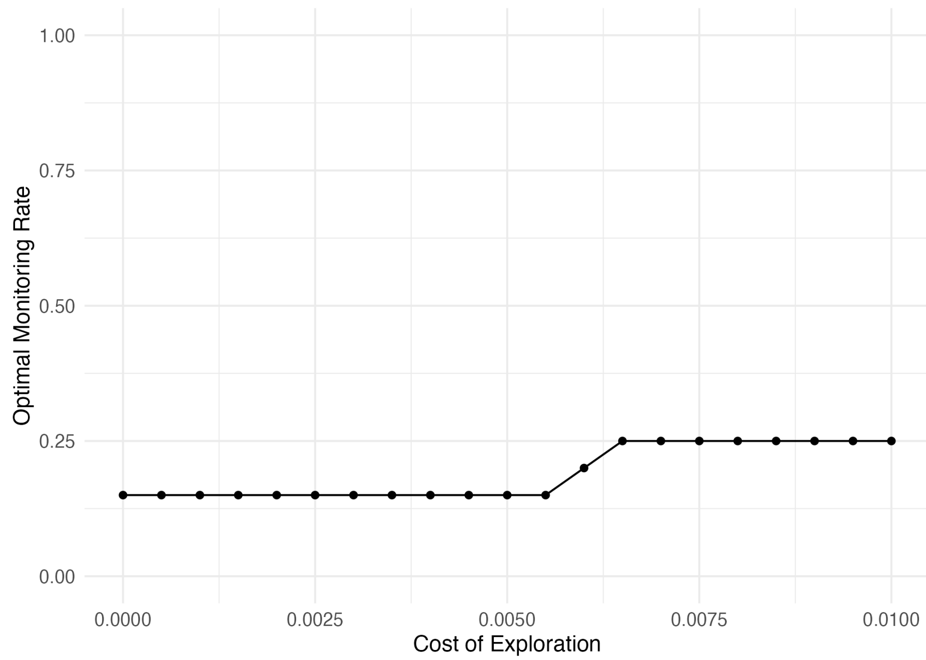

that is likely to have greater demand is more beneficial to a smaller firm.

Figure 4: Simulation Results - Arrival Ratio

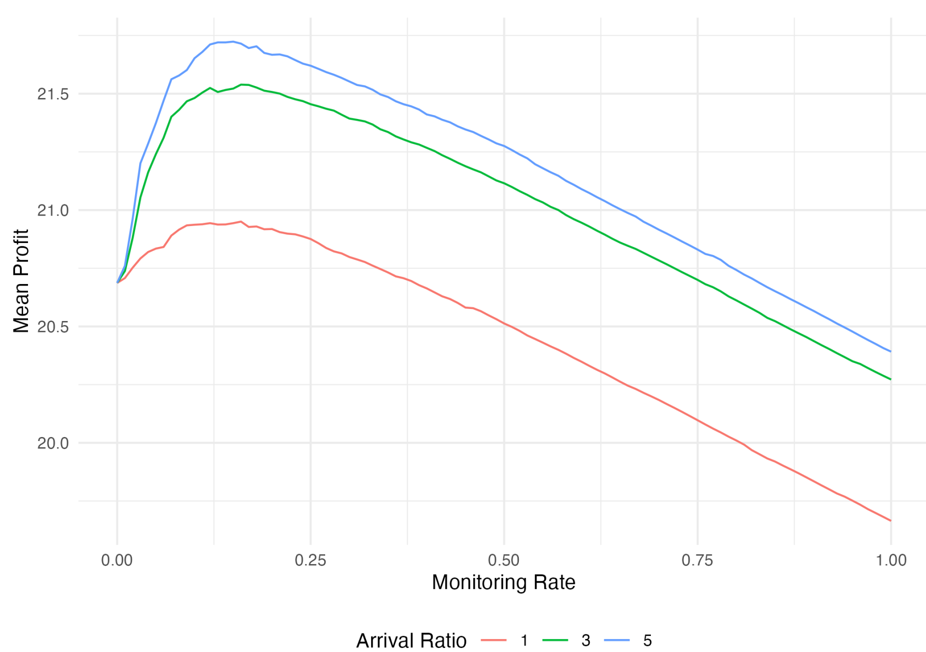

4 Data

This section presents institutional details about our empirical retail setting and data. We use these data to motivate the assumptions of our bandit model. And, we demonstrate a premise of the bandit model, which is that monitoring is associated with price adjustments – even when a retailer does not observe a competitor's demand.

4.1 Institutional Details and Available Data

The data was provided by a large online retailer that operates primarily in North America. The retailer offers millions of products, and employs pricing algorithms to price these products. Given the breadth of the retailer's assortment, we focus our study on a subset of categories that includes small appliances, tools and accessories. Due to confidentiality restrictions, we cannot provide more details but all of these items are commonly used throughout a home and are widely available online. Our data set includes 113,024 products; to put the size of our dataset in perspective, a large supermarket typically offers 40,000 to 60,000 products, and a Walmart Supercenter offers roughly 140,000 products.[3] Our product level information includes category, sub-category, brand, manufacturer, and manufacturer suggested retail price (MSRP). Our time-series spans a 6 month period between June 1 and December 9, 2019, and our final dataset includes 17,274,677 product–day observations.

For each product, we observe its daily price and an index for the units sold of that product at the focal website. The index disguises the actual demand for each product, but is perfectly correlated with actual demand. To aid with interpretation, we scale the index so that a zero value corresponds to zero demand, and a positive index value corresponds to positive demand. Revenue is calculated as price multiplied by the demand index and the products in our study contributed millions of dollars in revenue.[4] We do not observe customer shipping costs and delivery times (e.g., one- or two-day delivery, etc) as they vary by transaction characteristics (e.g., time of day, geographic location, minimum spend), but

[3] See https://corporate.walmart.com/newsroom/2005/01/06/our-retail-divisions.

[4] We further scale the revenue by another scalar to preserve confidentiality.

most of the orders for this retailer include free shipping.

For each product-day observation, we know whether the retailer monitored a competing retail website that sells the same product. When a competing website is monitored, we observe the price. Nearly all of the brands in our sample are available at other e-commerce retailers, which implies that it is possible to find products on competing websites. When a product is not monitored, the main reason is that the focal retailer's algorithm did not capture a competitor's price that day. We will later show that products differ in their frequency of competitive monitoring.

One might expect the retailer to monitor prices more frequently than once a day. But in our data, less than 1% of products have monitored prices from the same competitor within the same day.[5] A daily monitoring frequency is consistent with our finding that prices are sticky at both the focal firm and competing firms. When multiple competitor prices are gathered on the same day from a single competitor, they are typically identical. And roughly 90% of the product-day observations represent no change in price relative to the previous day for the focal firm or a competing firm. This is consistent with previous research that used daily price scraping (Cavallo and Rigobon 2016, Cavallo 2018) and showed infrequent price changes across retailers and product categories.

We recognize that other researchers document frequent price changes, particularly for more frequently purchased items such as grocery and convenience store products. For example, Brown and MacKay (2023) document high frequency pricing among popular branded allergy medicines. Aparicio et al. (2021) and Aparicio et al. (2024) document frequent price changes among some grocery products. We caution that these studies may differ from ours in several important ways. First, these papers consider frequently purchased goods (e.g., grocery) versus durable goods. Second, these papers compare pricing for a small number of products whereas we systematically study over 100,000 products. Third, these studies scrape prices and assume that competing retailers have access to this information. In contrast, we

[5]A few observations in our data are of a competitor that is monitored more than once per day (82,747). That is 1.8% of total monitored competitor observations, and less than 0.5% of the total product-day observations. In case of multiple prices we use the lowest price.

focus on what the focal retailer knows about competitive prices. As we will show, the focal retailer does not have full information about competing prices.

Overall, it seems that only a few (but large) sophisticated retailers in the market monitor and change prices frequently within the same day, while for most other retailers, a lower frequency of price monitoring and price changes (i.e., daily, not hourly) is the norm. In this sense, we view our data and observations as representative of many online, durable good retailers.

4.2 Summary Statistics

Because we are focused on the long-tail of products, we split the data to "popular" products and "long-tail" products. We use the Pareto rule to identify those top products that represent 80% of revenue and label these popular products. To classify products into popular and long-tail, we utilize product sales from the time the product was introduced on the platform until the day *before* the period of our data, which is May 31, 2019. Importantly, sales from June to December 2019 are not considered for classification as popular or long-tail. In our data, popular products are 0.9% of the total products, which represent 5.0% of products with positive sales during the time period of the data. Among the long-tail products (i.e., "unpopular") we identify two important sub-groups: *Group* 1 includes long-tail products that have both competitive monitoring and price changes by the focal firm; *Group* 2 includes long-tail products that have no competitive monitoring but do have price changes. There are additional long-tail products (3.4% of total products) that have no price changes throughout the entire time period. We omit these products from our analyses as they do not change our core findings.

In Column (1) of Table 2 we present summary statistics for our data. The average price of items in the data is $61.2, the average product-day demand index is 2.9 and the average product-day revenue index is 15.9. We also note that there is considerable variation in these metrics. On average, 41.5% of products are monitored at least once in our sample. If we consider monitoring at the product-day level, only 24.3% of observations involve monitoring.

On average, 0.9 competitors are monitored for a product and this metric includes zeros for nearly 60% of the products that are not monitored. Among those products that are monitored, the average number of competitors monitored in the six month time period of our sample is 2.3, whereas the maximum number of competitors monitored for a particular product is 9. Most products are widely available at other online sellers, as less than 1% are from brands exclusive to the focal retailer.[6] Only 18.0% of products have positive sales over the sample period. This is consistent with the online retailer offering a long-tail of product variety.

We observe that daily price changes are not frequent: only 10.8% of total observations represent any price change relative to the previous day for a particular product. To help understand whether these are large or small price changes, we classify price changes based on whether the change is at least 1% or at least 5% in absolute value. For example, the overall price change frequency is 10.8%, but if we only look at price changes of at least 1%, the frequency drops to 9.6%. The difference of 1.2% represents small price changes under 1%. If we consider price changes of more than 5%, which is an empirical threshold that some researchers have used to indicate a price promotion (e.g., Hitsch et al. 2021), the frequency of price changes drops to 5.5%.

The focal retailer monitors prices at fifteen large retailers and thirteen of these had annual sales over one billion dollars during the time period of our data. But monitoring is concentrated on two competitors; we refer to the largest as Competitor #1 and the next largest as Competitor #2. These are also the two largest competitors. When we consider all products that are monitored, 86% have at least one day of monitoring at Competitor #1 and 67% have at least one day of monitoring at Competitor #2. If we combine the thirteen other competitors, we find that 43% of monitored products are monitored at least once at any of these competitors. Therefore, our focus on the top competitors captures the bulk of

[6] A search of manufacturers' websites revealed the names of authorized e-commerce retailers. Using this information and manual searches, we found that the most popular branded products are offered on over 30 competing websites. Despite this fact, the retailer only monitors 15 retailers, consistent with theories of costly monitoring.

Table 2: Product-level summary statistics

	All Products		**Popular**		**Group 1**		**Group 2**	
	(1)		**(2)**		**(3)**		**(4)**	
	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.
Price	61.17	184.26	243.17	347.73	66.57	207.69	48.68	146.67
Demand Index	2.92	98.62	69.46	349.28	3.91	131.43	0.54	16.72
Revenue Index	15.87	459.73	635.76	3009.4	16.7	472.2	1.93	97.43
Any Monitoring? (%)	41.5%	0.49	82.1%	0.38	100.0%	0	0%	0
Monitoring (%)	24.3%	0.43	80.3%	0.4	47.0%	0.5	0%	0
#Competitors	0.94	1.39	3.15	2.41	2.24	1.25	0	0
Is Exclusive? (%)	0.9%	0.1	0.0%	0	0.0%	0	1.5%	0.12
Any Sales? (%)	18.0%	0.38	88.5%	0.32	33.1%	0.47	5.8%	0.23
Any Price Changes? (%)	10.8%	0.31	12.8%	0.33	12.2%	0.33	9.8%	0.30
1% Price Changes? (%)	9.6%	0.29	7.8%	0.27	10.3%	0.30	9.5%	0.29
5% Price Changes? (%)	5.5%	0.23	4.7%	0.21	5.7%	0.23	5.5%	0.23
# Unique Products	113,024		1,022		44,675		63,508	

Notes: *Group*1 and *Group*2 are all long-tail products. The 3,819 products that are included in Column (1) but do not appear in Columns (2)–(4) are products in the long-tail with no price changes during the entire time period. Monitoring (%) is at the product x day level. #Competitors is number of competitors monitored.

competitive monitoring. We observe similar magnitude of price changes among these top two competitors, as detailed in Web Appendix Table W1.

We also find that the retailer is roughly equally likely to increase or decrease price, which is consistent with price exploration; the average magnitude of price change if not splitting the data to positive and negative changes is 0.1%. Price increases are modest, with the median at 5% and the average price increase has a magnitude of 8.7%. Similarly, price decreases are modest, with the median at 5%, and the average price increase has a magnitude of 6.5%.[7]

Columns (2)-(4) of Table 2 highlight the differences between the different groups of products. The group of popular products have a higher price on average, higher demand and revenue indexes, and 80.3% of products have monitoring, with a high average monitoring

[7] Additional details about price changes are available in Web Appendix B.

frequency of 82.1% of product-days. By construction, all products in this group have sales. Compared to the popular products, the long-tail of products are lower priced, lower demand and lower revenue. In terms of monitoring, *Group* 1 long-tail products are monitored less than half of product-days on average (i.e., 47%), despite the fact that all of the products in the group have at least some monitoring. *Group* 2 long-tail products tend to have lower demand and revenues, but also are more likely to include exclusive brands, which may explain why some of the products in the group are not monitored. We note that we control for observable differences, like exclusivity, in our later analyses of *Group*1 and *Group*2 products. In terms of price changes, When we consider price changes among popular versus long-tail products, we observe 12.8% for the popular group, 12.2% for *Group* 1 and 9.8% for *Group* 2. Thus, *Group* 1 long-tail products exhibit more frequent price changes relative to *Group* 2 products, which do not include monitoring. However, when we examine price changes of at least 5% the frequency of price changes among the three groups is more similar. This suggests that the overall differences in price changes are due to small price changes among popular and long-tail products.

4.3 Monitoring Policy

Next, we examine the focal retailer's monitoring policy. While we observe detailed, daily data on monitoring of competitor prices by product, we do not know the specific monitoring policy of the focal retailer. Therefore, we infer the monitoring policy from our data. After extensive analysis, we found that monitoring rates were relatively stable among weeks. For example, if a product was monitored on a single day one week, then it was very likely that the product would also be monitored on a single day the following week.

To illustrate this pattern in the data, we focus on the 41.5% of products that were ever monitored. For each product, we calculate the number of days the product was monitored each week. Then, we compute the mode of this metric for each product.

In Table 3, we report the proportion of products for each mode and observe two types of policies. The first monitoring policy is a high frequency policy (mode=7) of monitoring every day. The second policy is a low frequency policy of monitoring 3 or fewer days (mode = 0,1,2,3). We note that monitoring 4, 5 or 6 days per week is very uncommon (less than 5% of products). Among popular products, 6.8% are monitored at a low frequency, 3.8% at a medium frequency, and 89.5% at a high frequency. Among the *Group* 1 products, 66.8% are monitored at a low frequency, 4.7% at a medium frequency, and 28.5% at a high frequency. Importantly, 89.2% of the 44,675 products in this groups are monitored at a mode of at least once per week. All of *Group* 2 products have a modal monitoring of zero by construction.[8] Overall, these monitoring frequencies are consistent with theories of costly monitoring, which is also a key feature in our bandit model.

Table 3: Modal Monitoring Policies

Modal Days/ Week Monitored	0	1	2	3	4	5	6	7
All Products	10.8%	21.4%	18.6%	13.9%	3.4%	0.8%	0.5%	30.6%
Popular	3.2%	0.6%	1.2%	1.8%	1.1%	0.4%	2.3%	89.5%
Group 1	10.8%	22.1%	19.4%	14.5%	3.5%	0.8%	0.4%	28.5%
Group 2	100%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%

Note: The frequencies reported in the table reflect the percentage of products that have modal days/week monitored within each group.

4.4 Monitoring as a Source of Information

Finally, we provide descriptive evidence that links competitive monitoring to price adjustments in the long-tail. We focus on whether the focal retailer changed its price after monitoring the price at Competitor #1. We use the monitoring data to determine whether Competitor #1 did not change price, increased price or decreased price. We then examine how the focal retailer price changed.

[8]In Web Appendix B.2 we examine these modes further and show that we rarely see a low frequency product monitored 7 days per week. Vice versa, we rarely see a high frequency product monitored at a low frequency.

We find that for products in the long-tail for which the retailer monitors Competitor #1, pricing behavior is quite stochastic. First, the retailer engages in price exploration even when they observe no price changes by the competitor. The focal retailer changes the price within the week of monitoring 42.7% of the time after observing no competitor price change; the focal retailer is slightly more likely to increase prices (22.3%) relative to decreasing prices (20.4%). When the competitor is observed to change the price, the retailer is more likely to explore prices in the direction of the competing retailer's price change. (For detailed information see Web Appendix C.1.) This observation motivated our operationalization of exploration in the bandit model.

Next, we assess how quickly the focal retailer adjusts prices within the week after they monitor the competitor. Table 4 displays the proportion of cases in which response occurs for each day. We find that after the competitor changed the price, most of the price changes occur in the first two days (63.3% for price decreases, and 58.6% for price increases). For comparison, we also examine the retailer's response after observing no price changes. In this case, the response is higher on day 1 (28.8%), which is substantially lower than cases in which a price change was observed (43.6%).

Table 4: Price change timing after monitoring

Day of change/ Competitor #1 Behavior	1	2	3	4	5	6	7
Any change	43.6%	17.9%	11.4%	8.9%	8.5%	5.8%	3.8%
Decrease	45.9%	17.4%	11.2%	8.6%	8.1%	5.4%	3.4%
Increase	39.8%	18.8%	11.8%	9.5%	9.2%	6.4%	4.4%
No change	28.8%	18.3%	14.9%	12.2%	10.3%	8.4%	7.1%

Note: The frequencies reported in the table reflect the percentage of cases with a price change at that day within each group. Each row sums up to 100%.

To further examine whether competitive monitoring is associated with price changes reaction by the focal retailer, we present in Web Appendix C.2 an "event study" type analysis that aims to isolate further how the focal retailer adjusts price in response to competitive price information. In that analysis, we only consider products for which the focal retailer hadn't yet had any sales and that have no price changes by Competitor #1 the prior four weeks (28 days). Then, we examine the focal retailer's price changes after observing a price change by Competitor #1. Similarly to the above, we find that most of the reaction to a price change seems to be within the first two days of observing a price change by the competitor, and persists to up to 6 days after. We also find that the focal retailer responds in the direction of the price change, and that the magnitude of the price change is proportional to the magnitude of price change by the focal retailer. Again, this indicates that the focal retailer seems to respond to information obtained via competitive monitoring.

In sum, when there are no price changes by the largest competitor, the focal retailer engages in price exploration for long-tail products. When the focal retailer observes a price increase (decrease) by the largest competitor, this is probabilistically associated with an upward (downward) price adjustment and exploration of higher (lower) prices. Finally, price adjustments typically occur within the first few days of observing a change in the competitor's price.

4.5 Summary

In this section, we demonstrated how the empirical data motivated several key assumptions in our bandit model.

- Costly monitoring: we demonstrated two pieces of evidence consistent with costly monitoring. The first is that the firm monitors only a subset of 15 competitors, and out of them, monitoring is especially concentrated on the two largest competitors. The second is that the intensity of monitoring is heterogeneous: certain products are subject to daily monitoring, others to weekly, and some are never monitored.

- Single large competitor: because 86% of the monitored products were monitored at Competitor #1, our bandit model focuses on a single competitor as well.

- Prices are sticky: the firm typically charges a regular price and price changes only occur 11% of the time.

- Sources of information utilized: the firm engages in both price exploration and competitor price monitoring activities.

- Monitoring as a source of information about prices: monitoring a competitor is associated with subsequent price changes at the focal retailer.

5 Empirical Evidence on Competitive Monitoring

In this section, we show consistency between the results of our model and the empirical context. Specifically, we show that (i) a higher level of monitoring competitors' prices in the long-tail is associated with higher revenues; (ii) monitoring and price exploration are substitutes; and (iii) monitoring a larger competitor is associated with larger revenues.

Our analyses is organized into three subsections. First, we examine the relationship between past monitoring and future revenues (Section 5.2). Second, we analyze the relationship between past monitoring, price exploration, and future revenues (Section 5.3). Third, we examine the relationship between monitoring the largest competitor and future revenues of the focal firm (Section 5.4).

5.1 Overview of Analyses

In the analyses in this section, we seek to understand how competitive monitoring and price changes (i.e., price exploration) are associated with future sales. The core idea in our subsequent analyses is that price adjustments should eventually lead to an increase in demand and/or revenue. In the bandit model in Section 3, sales also occur without any monitoring, but at a lower price, and the benefit of monitoring is gaining higher revenues and margin. In our empirical data, the long-tail products have had zero sales since inception. In this case, if monitoring competitive prices is informative, this should also be associated with both a change in price and a change in unit sales. For consistency with the bandit model analyses, and because we do not observe costs, we focus on revenues in this section, although our results also hold for unit sales (see Web Appendix D).

Our empirical approach in this section compares products across the two long-tail groups (*Group* 1 and *Group* 2) in two time periods to examine the differences between these groups over time. We classify the first three months of our data as the pre-period and the last three months of our data as the post-period. To construct the sample, we consider products from *Group* 1 and *Group* 2 that have zero sales in the pre-period. Moreover, we ensure that these products never had any sales even before this period (recall from section 4.2 that we also observe whether product had any sales before the data sample time period). We also confirm that the *Group* 1 products had monitoring in both pre- and post- time periods. This yields a total of 19,228 *Group* 1 products and 26,986 *Group* 2 products and two time periods.

We collapse the data to two time periods because all of our products have zero demand in the first time period and demand is sparse in the second time period. This helps alleviate concerns with serial correlation (Bertrand et al. 2004). However, we run additional robustness checks and show that our results are robust to using daily level data rather than two periods (see Web Appendix D.2). Finally, we perfomed a robustness check with a two-month post-period to remove any concerns that holiday transactions in November and early December were driving our results (see Web Appendix D.3).

Recall that in our setup, *Group* 1 products have monitoring and *Group* 2 products have no monitoring; both groups of products have price changes. Our analyses compares sales outcomes for the two groups, but a potential concern is that the products in the two groups are different leading to different monitoring policies. We address this concern in two ways. First, our design matches on pre-period demand of zero (i.e., no sales prior to September 1st, 2019). Second, we use generalized full matching (Sävje et al. 2021) to (partially) match on observables (see Web Appendix E for details.). While this approach matches on observables, it does not match on unobservables, which is a limitation discussed later. Therefore, we see this analyses as suggestive evidence of correlations in the data and do not claim to establish a clear causal link.

5.2 Monitoring and Future Revenues

In this subsection, we examine the correlation between monitoring in the pre-period and demand in the post-period. We estimate the following OLS regression:

$$Y_{it} = \alpha Post_t + \beta Post_t \times Monitor_i + \eta X_{it} + \zeta_i \quad (1)$$

where Y_{it} is a revenue measure for product i at time t, $Post_t$ indicates the last three months of the data, $Monitor_i$ indicates whether a product has monitoring (belongs to *Group* 1), and ζ_i are product fixed effects. The term X_{it} includes product-time controls: frequency of price changes, the variation in prices measured as the coefficient of price variation (the standard deviation of prices over the mean of prices $\frac{\sigma}{\mu}$), and the average price during the period.

This specification leverages the differences in monitoring among *Group* 1 and *Group* 2 products. The coefficient α captures the increase in revenues from zero to a positive value for all products. The coefficient β captures the differential increase in revenues for *Group* 1 products. To further decompose the impact of monitoring among *Group* 1 products, we also estimate specifications with *Post* interacted with $Med\ Monitor_i$ (i.e., medium frequency of monitoring) and $High\ Monitor_i$ (i.e., high frequency of monitoring).

Table 5 presents the results from six different specifications. Column 1 presents the results with no fixed effects, Column 2 contains product fixed effects, Column 3 uses a matched sample and includes both product fixed effects and weighs observations based on the matched weights.[9] Not surprisingly, the *Post* coefficient is positive and significantly different than zero, as there are no sales at all in the "pre-period." Across all three columns, we find that those products that had monitoring (i.e., belong to *Group*1) have higher revenues in the post-period. The increase in the revenue index is 24.9 units in the model with product fixed effects and matching (Column 3); this is relative to average daily revenue of 16.7 in *Group* 1. In Column 3 the *Post* coefficient is 28.5, which shows that ceteris paribus *Group* 2

[9]Note that in the matched regression, about 30% of the observations have weight below 0.0001, corresponding to those observations in *Group* 2 that do not have common support on observables with *Group* 1.

products have average revenues of 28.5 and *Group* 1 products have revenues of 53.5. Thus, *Group* 1 products have roughly 85% higher revenues relative to *Group* 2 products.

Table 5: Association of Monitoring and Future Revenues

	(1)	(2)	(3)	(4)	(5)	(6)
Post	7.769***	11.615***	28.531***	12.284***	15.850***	26.227***
	(2.427)	(2.286)	(3.548)	(2.122)	(1.968)	(2.599)
Post X Monitor	40.933***	43.051***	24.936***			
	(3.017)	(4.952)	(5.812)			
Post X Med Monitor				57.651***	56.507***	46.100**
				(11.651)	(17.820)	(17.901)
Post X High Monitor				288.550***	285.835***	275.538***
				(7.476)	(44.744)	(45.055)
Mean Price	0.028***	-1.306***	-0.950**	0.027***	-1.052***	-0.821**
	(0.005)	(0.422)	(0.399)	(0.005)	(0.388)	(0.360)
Price Change Frequency	10.844	44.064	33.179	3.837	35.927	47.469
	(19.050)	(54.762)	(41.762)	(18.888)	(51.922)	(40.604)
Price Variation	10.634	-60.175	-56.579	33.086	0.580	-20.746
	(21.738)	(42.355)	(39.644)	(21.404)	(38.542)	(37.279)
# Obs	92,428	92,428	92,428	92,428	92,428	92,428
Adjusted R2	0.004	0.005	0.005	0.018	0.018	0.017
Product FE	No	Yes	Yes	No	Yes	Yes
Matching	No	No	Yes	No	No	Yes

* significant at 10%; ** significant at 5%; *** significant at 1% level.

In Columns 4 through 6 in Table 5, we replace the *monitor* indicator with two indicator variables that indicate the product monitoring policy: low (omitted category, includes modal monitoring days are 0, 1, 2, 3), medium (modal monitoring days are 4, 5, 6), and high (modal monitoring is 7 days). Overall, we find a monotone relationship, where higher monitoring frequency is correlated with higher revenues. Importantly, this is consistent with our simulations in the bandit model where we showed that increased monitoring led to increased revenue (see Table 1).[10]

[10] In Web Appendix D we repeat the analysis in Table 5 using the sales index instead of the revenue index and find similar results; monitoring is positively correlated with unit sales in the post-period.

5.3 Monitoring, Price Changes, and Future Revenues

We now extend the analysis to explore how past price changes are associated with future revenues. Recall that there is extensive price exploration for both *Group* 1 and *Group* 2 products. To measure these associations, we estimate the following model:

$$\begin{aligned} Y_{it} = \alpha Post_t + \beta Post_t \times Monitor_i + \gamma Post_t \times HighFreq_i + \\ \delta Post_t \times Monitor_i \times HighFreq_i + \eta X_{it} + \zeta_i \end{aligned} \tag{2}$$

where the new variable $HighFreq_i$ indicates whether the frequency of price changes for a product was greater than the median in the *pre-period*. Because all of our data includes price changes, we median split the data to indicate more frequent price changes which we use as a proxy for more intense experimentation. We then examine the relative magnitude and sign of β, γ and δ, which allow us to decompose the association of monitoring and frequent price changes with revenues. Notice that we include interactions of *Post* with monitoring, price change frequency, and a three-way interaction of all variables (in models without product fixed effects, we also include the interaction between monitoring and high frequency of price changes). If price changes and monitoring are complements with respect to associations with demand then we expect $\delta > 0$. But if price changes and monitoring are substitutes then we expect $\delta < 0$.

Table 6 reports the results of four different model specifications. In Columns 1 and 2 of the table we use each observation with equal weights, adding product FE in Column 2. In Columns 3 and 4 we use the weights based on matching. We find that monitoring has a positive and significant coefficient in all four models. We also find that those products that had more frequent price changes (proxy for price exploration) in the pre-period have higher revenues in the post-period. Since we specify both variables as binary, we can qualitatively compare the magnitude of monitoring versus price changes. We observe in Column 4 that the association between monitoring and indexed revenues is roughly 67 while more frequent price changes yields an association of roughly 21. This demonstrates that past monitoring has a stronger association with future revenues than past high frequency price changes (i.e.,

Table 6: Association of Monitoring, Price Changes and Future Revenues

	(1)	(2)	(3)	(4)
Post	3.058	9.985***	14.588***	17.837***
	(2.969)	(3.396)	(3.471)	(3.872)
Post X Monitor	78.718***	79.723***	68.170***	68.994***
	(4.811)	(10.733)	(5.288)	(11.052)
Post X High Frequency	12.566***	9.607**	21.397***	20.809***
	(4.091)	(4.058)	(4.085)	(6.678)
Post X Monitor X High Freq	-59.441***	-59.533***	-69.292***	-68.476***
	(7.329)	(12.242)	(7.505)	(13.397)
Mean Price	0.026***	-1.299***	0.021***	-0.943**
	(0.005)	(0.418)	(0.005)	(0.393)
Price Change Frequency	20.000	-20.846	15.239	15.113
	(20.743)	(63.244)	(17.030)	(46.113)
Price Variation	16.842	-46.236	15.996	-46.418
	(21.746)	(41.930)	(19.696)	(39.342)
Monitor X High Frequency	-1.176		-0.264	
	(3.496)		(3.487)	
# Obs	92,428	92,428	92,428	92,428
Adjusted R2	0.005	0.006	0.005	0.006
Product FE	No	Yes	No	Yes
Matching	No	No	Yes	Yes

* significant at 10%; ** significant at 5%; *** significant at 1% level.

price exploration).

Additionally, the coefficient of the three-way interaction between post, monitor, and high frequency in Table 6 is negative and of similar in magnitude as the $Post \times Monitor$ interaction. This shows that products with both monitoring and high frequency of price changes in the post-period don't reap additive benefits. And, we obtain similar conclusions when using the sales index as the outcome variable (available from authors). Together, this suggests that past monitoring and past price changes should be viewed as substitutes, and both are associated future demand and revenues. This is consistent with the bandit model where price exploration and monitoring are substitutes with respect to learning about demand in the long-tail.

5.4 Monitoring a Larger Competitor and Future Sales

The idea that a retailer can learn from a competitor assumes that the other retailer is exploring price and eventually generating demand. Demand for long-tail products is highly unpredictable, but presumably a larger competitor is more likely to observe positive demand. In this subsection, we investigate this possiblity. We add another variable to our regressions that captures the incremental effect of monitoring the largest competitor — Competitor #1. We repeat the same analysis but include an interaction term for $Post \times MonitorCompetitor\#1$.

Table 7: Monitoring Competitor #1

	(1)	(2)	(3)	(4)
Post	3.057	9.917***	14.591***	17.793***
	(2.969)	(3.388)	(3.471)	(3.869)
Post X Monitor	63.574***	63.435***	52.959***	52.649***
	(6.754)	(10.695)	(7.272)	(11.050)
Post X Monitor Comp #1	19.478***	20.935**	19.550***	21.019**
	(6.098)	(9.765)	(6.415)	(9.755)
Post X High Frequency	12.568***	9.669**	21.396***	20.870***
	(4.091)	(4.053)	(4.085)	(6.676)
Post X Monitor X High Freq	-60.902***	-60.966***	-70.755***	-69.986***
	(7.342)	(12.399)	(7.520)	(13.549)
Mean Price	0.026***	-1.299***	0.021***	-0.942**
	(0.005)	(0.418)	(0.005)	(0.393)
Price Change Frequency	20.166	-17.722	15.321	17.072
	(20.742)	(62.990)	(17.029)	(45.970)
Price Variation	16.081	-51.139	15.427	-50.127
	(21.746)	(42.616)	(19.696)	(39.788)
Monitor X High Frequency	-1.180		-0.268	
	(3.496)		(3.487)	
# Obs	92,428	92,428	92,428	92,428
Adjusted R2	0.005	0.006	0.005	0.006
Product FE	No	Yes	No	Yes
Matching	No	No	Yes	Yes

* significant at 10%; ** significant at 5%; *** significant at 1% level.

The results of four different specifications are in Table 7. When we compare the coefficients in Table 6 and Table 7 the main differences are the $Post \times Monitor$ and $Post \times Monitor\ Competitor$ #1 interactions. Adding the interaction with Competitor #1 allows us to decompose the monitoring coefficient and demonstrate the relative importance of monitoring a large competitor. If we use the model with fixed effects and matching (Column 4), the total monitoring effect is $52.649 + 21.019 = 73.668$. This suggests that roughly a third of the relationship between monitoring and revenues is through monitoring Competitor #1. Overall, we find that all past monitoring is associated with future revenues, but monitoring the largest competitor has a stronger association. This is consistent with the comparative statics in the bandit model where we showed that profit of the smaller firm increases as the competitor generates more demand.

In Web Appendix D.3, we test the robustness of our results to shortening the "post" period to two months, September and October, excluding the month of November and early December that is typically associated with higher seasonal sales. While other results hold, the coefficient of the interaction of monitoring Competitor #1 is positive but no longer significant (see Table W10). This is likely due to the lower number of sales during the new "post" period. Our interpretation of this result is that it is particularly important to monitor competitors in the long-tail in periods of high demand, as the signal about demand is likely more informative at that time.

5.5 Summary

In this section, we showed consistency between key results from our model and the empirical data. First, in Table 5, we showed that the intensity of monitoring is correlated with higher revenues. A medium and a high frequency of monitoring was associated with $46 and $275 additional revenues relative to a low frequency baseline, respectively. This is consistent with the results in Column (6) of Table 1, where revenues are increasing monotonically at the monitoring rate.

Second, in Table 6, we showed that monitoring and price exploration are substitutes. When the firm engages in high-frequency price changes and monitoring, the combined association with prices is not additive. This is consistent with the results in Column (1) of Table 1, where price exploration decreases as monitoring increases.

Third, we showed in Table 7 that monitoring the largest competitor is an important part of the monitoring benefit, as it accounts for a significant portion of the monitoring coefficient. This is consistent both with the Table 1 that shows that monitoring a large competitor is associated with higher revenues than not monitoring, and with Figure 4 that shows that a competitor with higher arrival rates relative to a focal firm yields better outcomes.

6 Discussion and Conclusion

While the concept of the long-tail products has existed since the emergence of e-commerce, very little is known about how retailers adjust price to generate demand for these products. In this paper, we develop a bandit model for long-tail products where demand signals are rare. Our algorithm combines competitor prices and price exploration to adjust prices of long-tail products.

A core contribution of our paper is to highlight that monitoring a large competitor is an important source of information for long-tail products. This is a key prediction of our bandit model, and empirically we show that past monitoring of the largest competitor's prices is associated with a substantial increase in future revenues.

We also show that monitoring and price exploration are substitutes. As a consequence, when price exploration is more expensive a firm will engage in more price monitoring. Alternatively, when monitoring is more costly a firm may engage in more price exploration.

To simplify the exposition of the bandit model, our simulations considered a case where prices gradually rise as retailers learn about demand. This case conveys the main intuitions from the bandit model in the simplest possible context. One can extend this model to allow for situations where price may either increase or decrease over time as a retailer learns about demand. For example, if marginal cost is zero and the distribution of valuations followed a normal distribution, then the optimal starting price may be the mean. Because the true value of V may be either greater or less than the mean, prices may either increase or decrease as the retailer learns.

A strength of our paper is that we demonstrate consistency between predictions of the bandit model and our empirical data. However, we recognize that our empirical evidence is based on observational data. While we utilize quasi-experimental methods to identify causal relationships, these come with limitations. Perhaps the greatest threat to giving our empirical results a causal interpretation is unobservable variables that are correlated with both the independent variables (e.g., monitoring rate) and the outcome (e.g., revenue), which leads to a confound. Given this, we prefer a more cautious interpretation of our findings (i.e., correlations and descriptive). To truly measure the causal impact of monitoring on demand, future work may have to randomly assign products to different monitoring policies and pricing policies. A carefully designed experiment would allow one to disentangle the relative contributions of monitoring and price exploration. We view our paper as a first step towards more fully understanding how retailers may learn about demand for long-tail products.

7 Funding and Competing Interests

All authors certify that they have no affiliations with or involvement in any organization or entity with any financial interest or non-financial interest in the subject matter or materials discussed in this manuscript. The authors have no funding to report. The data provider (who chose to remain anonymous) has a right to remove its intellectual property or trade secrets from the paper.

References

H. Adam, P. He, and F. Zheng. Machine learning for demand estimation in long tail markets. *Management Science*, 70(8):5040–5065, 2024.

C. Anderson. The long tail. *Wired (San Francisco, Calif.)*, 12(10):170, 2004. ISSN 1059-1028.

C. Anderson. *The long tail: Why the future of business is selling less of more*. Hyperion, 2006.

D. Aparicio and K. Misra. Artificial intelligence and pricing. In K. Sudhir and O. Toubia, editors, *Artificial intelligence in marketing*, pages 103–124. Emerald Publishing Limited, 2023.

D. Aparicio, D. Eckles, and M. Kumar. Algorithmic pricing and consumer sensitivity to price variability. Working paper, MIT, July 2021.

D. Aparicio, Z. Metzman, and R. Rigobon. The pricing strategies of online grocery retailers. *Quantitative Marketing and Economics*, 22(1):1–21, 2024.

H. Bar-Isaac, G. Caruana, and V. Cuñat. Search, design, and market structure. *American Economic Review*, 102(2):1140–60, April 2012.

M. Bertrand, E. Duflo, and S. Mullainathan. How much should we trust differences-in-differences estimates? *Quarterly Journal of Economics*, 119(1):249–275, 2004.

Z. Y. Brown and A. MacKay. Competition in pricing algorithms. *American Economic Journal: Microeconomics*, 15(2):109–156, 2023.

E. Brynjolfsson, Y. Hu, and D. Simester. Goodbye Pareto principle, hello long tail: The effect of search costs on the concentration of product sales. *Management Science*, 57(8): 1373–1386, 2011.

E. Calvano, G. Calzolari, V. Denicolò, and S. Pastorello. Artificial intelligence, algorithmic pricing, and collusion. *American Economic Review*, 110(10):3267–97, October 2020.

X. Cao and J. Zhang. Preference learning and demand forecast. *Marketing Science*, 40(1): 62–79, 2021.

A. Cavallo. Scraped data and sticky prices. *Review of Economics and Statistics*, 100(1): 105–119, 2018.

A. Cavallo and R. Rigobon. The billion prices project: Using online prices for measurement and research. *Journal of Economic Perspectives*, 30(2):151–78, May 2016.

L. Chen, A. Mislove, and C. Wilson. An empirical analysis of algorithmic pricing on amazon marketplace. In *Proceedings of the 25th International Conference on World Wide Web*, page 1339–1349, 2016.

J. Choi and D. R. Bell. Preference minorities and the internet. *Journal of Marketing Research*, 48(4):670–682, 2011.

M. Clayton. The long tail: Infinite consumer choice, September 2019. https://www.pocketbook.co.uk/blog/2019/09/03/long-tail-infinite-consumer-choice/.

J. Cooprider and S. Nassiri. Science of price experimentation at amazon. *Business Economics*, 58:34–41, 2023.

M. Fisher, S. Gallino, and J. Li. Competition-based dynamic pricing in online retailing: A methodology validated with field experiments. *Management Science*, 64(6):2496–2514, 2018.

D. Fleder and K. Hosanagar. Blockbuster culture's next rise or fall: The impact of recommender systems on sales diversity. *Management Science*, 55(5):697–712, 2009.

R. Ganti, M. Sustik, Q. Tran, and B. Seaman. Thompson sampling for dynamic pricing. *arXiv preprint arXiv:1802.03050*, 2018.

K. T. Hansen, K. Misra, and M. M. Pai. Frontiers: Algorithmic collusion: Supra-competitive prices via independent algorithms. *Marketing Science*, 40(1):1–12, 2021.

O. Hinz, J. Eckert, and B. Skiera. Drivers of the long tail phenomenon: An empirical analysis. *Journal of Management Information Systems*, 27(4):43–70, 2011.

G. J. Hitsch, A. Hortaçsu, and X. Lin. Prices and promotions in U.S. retail markets. *Quantitative Marketing and Economics*, 19:289–368, 2021.

Y. Huang, P. B. Ellickson, and M. J. Lovett. Learning to set prices. *Journal of Marketing Research*, 59(2):411–434, 2022.

J. Miklós-Thal and C. Tucker. Collusion by algorithm: Does better demand prediction facilitate coordination between sellers? *Management Science*, 65(4):1552–1561, 2019.

M. Mussi, G. Genalti, F. Trovò, A. Nuara, N. Gatti, and M. Restelli. Pricing the long tail by explainable product aggregation and monotonic bandits. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 3623–3633, 2022.

G. Oestreicher-Singer and A. Sundararajan. Recommendation networks and the long tail of electronic commerce. *MIS Quarterly*, 36(1):65–83, 2012.

T. W. Quan and K. R. Williams. Product variety, across-market demand heterogeneity, and the value of online retail. *The RAND Journal of Economics*, 49(4):877–913, 2018.

F. Sävje, J. S. Sekhon, and M. J. Higgins. quickmatch: Quick generalized full matching, 2018. https://CRAN.R-project.org/package=quickmatch.

F. Sävje, M. J. Higgins, and J. S. Sekhon. Generalized full matching. *Political Analysis*, 29 (4):423—447, 2021.

I. N. Weaver, V. Kumar, and L. Jain. Nonparametric pricing bandits leveraging informational externalities to learn the demand curve. Working paper, SSRN, May 2024.

H. Yang. Targeted search and the long tail effect. *The RAND Journal of Economics*, 44(4): 733–756, 2013.

Web Appendix

A Bandit Model Details

A.1 Deriving the profit maximizing price for the initial period

Because $V \sim U[\mu - x, \mu + x]$,

$$\text{Prob}(V < p) = \frac{p - (\mu - x)}{2x} \implies \text{Prob}(p \leq V) = 1 - \frac{p - (\mu - x)}{2x} = \frac{\mu + x - p}{2x}$$

Profits are then given by

$$\pi(p) = (p - c)\text{Prob}(p \leq V) = (p - c)\left(\frac{\mu + x - p}{2x}\right)$$

Solving for the optimal price

$$\begin{aligned} \frac{\partial \pi(p)}{\partial p} &= \frac{\mu + x - p}{2x} - \frac{p - c}{2x} \\ &= -\frac{p}{x} + \frac{\mu + x + c}{2x} = 0 \\ &\implies p^* = \frac{\mu + x + c}{2} \end{aligned}$$

The optimal initial focal price for $V \in [\mu - x, \mu + x]$ is then

$$p(1) = \begin{cases} \mu - x & p^* = \frac{\mu+x+c}{2} \leq \mu - x \\ p^* & p^* = \frac{\mu+x+c}{2} > \mu - x \end{cases}$$

A.1.1 Stopping rules details

(a) Cost vs value stopping rule There is no guarantee that the potential gains from exploration (i.e. $x + c - p^*(t)$) are larger or smaller than any costs from exploration (i.e. $c(\beta)$ and forgone revenue $\beta\alpha p^*(t)$). Any comparisons of expected return vs expected cost depend on the true value of V, which is unknown.

Experimentation stops when the incremental price gain from exploration is lower than the cost of exploration.

Specifically, we compare the potential *incremental* gain from exploration to the cost of exploration. The incremental gain at time t is: $(Avg\{support(V)\} - p^*(t)) \times Prob(sale)$, where $Prob(sale) = \alpha \times Prob(p^*(t) < V)$, therefore, we stop experimenting when: $(Avg\{support(V)\} - p^*(t)) \times \frac{\alpha}{2} < c(\beta)$.

(b) Final stopping rule If S periods have occurred without an update to the focal price, experimentation stops.

B Summary Statistics

B.1 Price Changes

Table W1 presents summary statistics for daily price changes for the focal retailer and for the top two competitors.

Panel B of Table W1 presents the overall price changes at the two largest monitored competitors. To calculate price changes among these two competitors, we compute changes relative to the last time the focal retailer monitored each competitor. We observe that competitors also change prices relatively infrequently (12%–13%) in this product category. While the average magnitude of price changes is slightly higher for the competitors relative to the focal retailer, this difference shrinks when we consider larger price changes.

Table W1: Summary statistics: Price changes

	Frequency					**Magnitude**	
	Any	**1%**	**5%**	**Positive**	**Negative**	**Positive**	**Negative**
	(1)	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**	**(7)**
Overall	10.8%	9.6%	5.5%	5.2%	5.6%	8.7%	-6.5%
Panel A: Product Groups							
Popular	12.8%	7.8%	4.7%	6.8%	6.0%	7.3%	-6.8%
Group 1	12.2%	10.3%	5.7%	6.1%	6.1%	9.3%	-6.6%
Group 2	9.8%	9.5%	5.5%	4.5%	5.3%	8.1%	-6.4%
Panel B: Major Monitored Competitors							
Competitor #1	12.5%	9.4%	5.9%	4.9%	7.6%	11.4%	-14.6%
Competitor #2	13.3%	11.4%	6.1%	5.3%	8.0%	4.4%	-5.5%

Note: Frequency indicates which fraction of the observations include price changes. The frequency of changes includes any price change compared to the previous price for that product (Column 1), a change of at least 1% (Column 2), and a change of at least 5% (Column 3). Column 4 indicates any positive price changes (price increase), and Column 5 indicates any negative price changes (price decrease). Column 6 and 7 present the average magnitude of positive and negative price changes, respectively.

B.2 Monitoring Policies Frequency

In this Appendix, we examine for each modal frequency the variation in number of days monitored. We use all the products that have monitoring and plot histograms of the weekly monitoring frequency by mode in Figure W1. If we consider mode = 7 (bottom right of the figure), we observe that 7 days is the most common frequency (by definition, at 74%) followed by 6 and 0 days. The density of observations on days 0 through 6 provides an indicator of the consistency of the 7 day monitoring policy. Since there is little mass on other days, it suggests that the high frequency policy has little variation.

For modes 4, 5 and 6, the histograms are very sparse, as these are not common policies. For low frequency policies (mode = 0, 1, 2, 3), there is more variation in the number of days monitored. For example, when the mode is 1 day per week, it is common to observe 0 or 2 days of monitoring. Thus, a product with a mode of once per week may occasionally change to zero or twice per week. When the mode is 3 days per week it is common to see 2 or four days per week. Most importantly for our characterization, we rarely see a low frequency product monitored 7 days per week. Vice versa, we rarely see a high frequency product monitored at a low frequency.

B.2.1 Monitoring Randomness within a Week

Next we explore whether monitoring is random within a particular week. Our analysis shows that there is some consistency in the rate of monitoring among weeks. But, within a week we observe that monitoring is much less consistent. If monitoring was perfectly random among days, then we would expect to observe 14.3% = 1/7 as the probabilities each day. To investigate this prediction, we focus on products that have a modal monitoring of one day per week. For each product, we calculate the proportion of times that monitoring occurs that day (e.g. Saturday). While there is variation in monitoring rates among products, the results in Figure W2 illustrate that on average, the proportion of times a specific product is monitored on a specific day is 1/7. Note that the Figure illustrates a box-plot, however because the 25th, 50th, and 75th percentile are all around 1/7, the box has no area. This

Figure W1: Distribution of total weekly monitoring, by mode

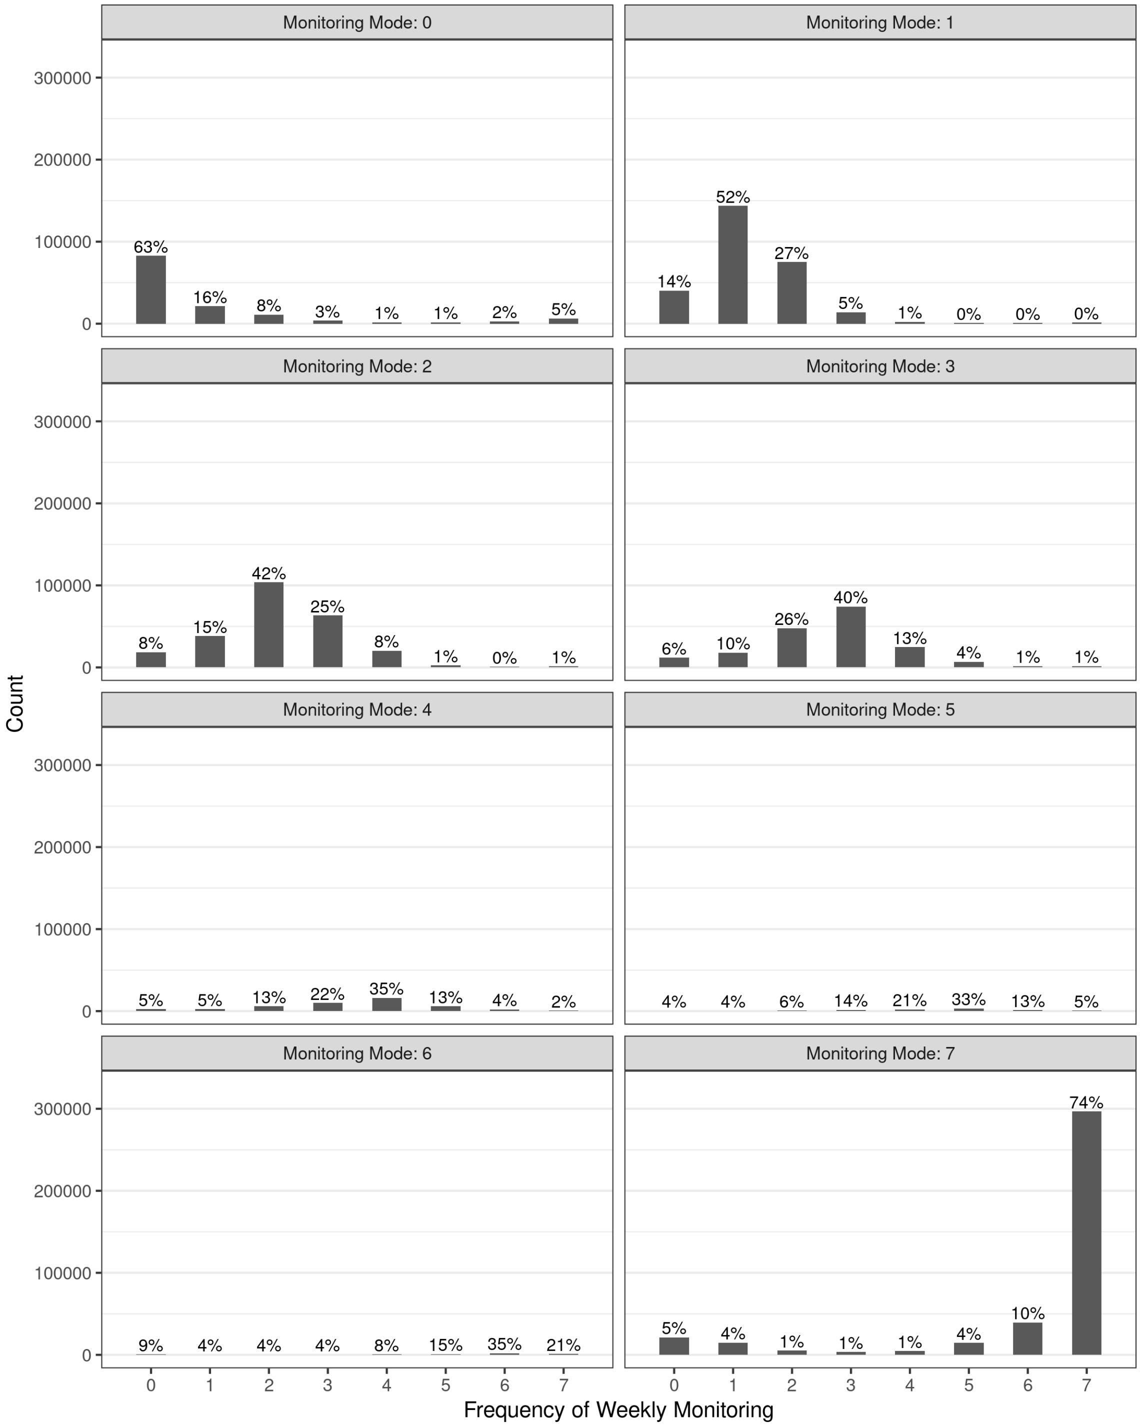

further demonstrates how monitoring is random across days of the week.

Figure W2: Monitoring Proportion for Each Weekday

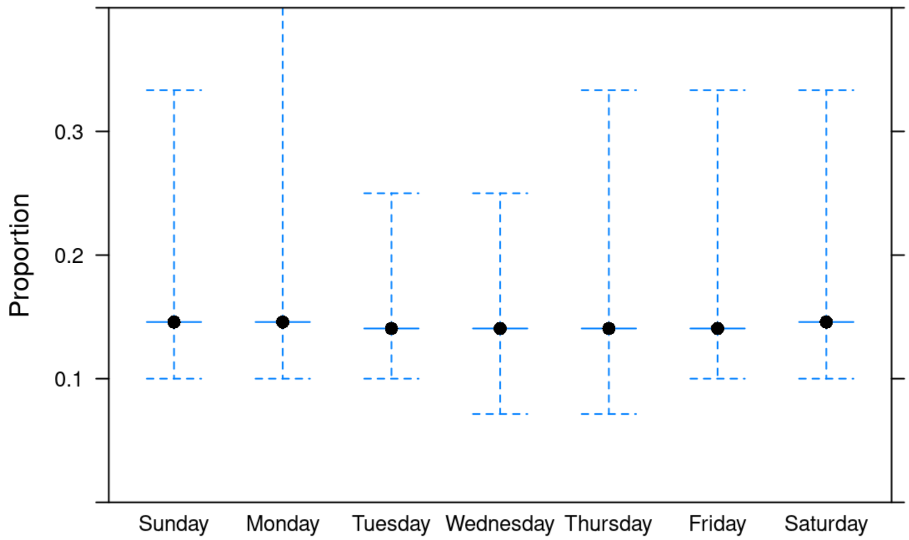

Note: This figure shows the proportion of monitoring on each weekday for all products with a modal monitoring of once per week

To further investigate the randomness among days, we calculate a transition matrix for products with modal monitoring of once per week. For a given weekday (e.g., Monday), we calculate the probability that each weekday is monitored the next week. We do this for products that have a modal monitoring of once per week[W1] and the results are summarized in Table W2.

Again, we observe that the unconditional probability of monitoring is around 14% for all days (Sunday and Monday have a slightly higher probability of 14.8%.). We observe that the

[W1] To compute this, we use all weeks for which the previous week had exactly one day of monitoring to create this matrix. If the previous week had one day of monitoring, and the next week has N days on monitoring, then each day this week with monitoring is assigned weights $\frac{1}{N}$. We note that this transition matrix is a bit noisy as there are weeks with zero monitoring as well as weeks with two or more days monitored. Indeed, if the previous week had no monitoring or the focal retailer monitored competitors on more than two days that, we omit that case from the analysis.

conditional probabilities are not uniform. Most notably, if monitoring starts on Thursday, then there is a 47.1% chance that Tuesday is monitored the next week, and if monitoring starts on a Friday, there is a 34.7% chance that Wednesday is monitored next week. Finally, examining the diagonal, except for starting on Sunday, which has the highest probability to stay on Sunday (29.6%), there is substantial transition from day to day the next week.

When we combine the evidence in Table W2 and Figure W2, it suggests that monitoring within the week is stochastic. While the distribution of conditional probabilities is not uniform among all days, monitoring is equally distributed across the week. As a consequence, it is unlikely that a competitor can behave strategically and send misleading price signals. Figure W3 further demonstrates that within each mode of monitoring, the day of the week in which monitoring occurs is fairly uniform.

We also explored in Figure W4 whether monitoring polices were related back to specific dates within a month (e.g., first day of the month, 15th of the month, end of month). We find no evidence that monitoring is concentrated on specific dates. Instead, a weekly policy with random monitoring among days best characterizes our data.

Table W2: Weekday Probability Transition Matrix monitoring

	Last Week							
This Week	Sunday	Monday	Tuesday	Wednesday	Thursday	Friday	Saturday	Unconditional
Sunday	29.6%	23.7%	30.2%	2.9%	1.4%	5.8%	8.1%	14.8%
Monday	14.0%	19.8%	15.6%	29.5%	4.1%	3.0%	6.1%	14.8%
Tuesday	13.7%	9.9%	11.8%	20.8%	47.1%	7.0%	4.6%	14.1%
Wednesday	6.1%	10.4%	5.8%	10.7%	19.4%	34.7%	5.3%	14.1%
Thursday	2.3%	4.4%	2.1%	4.2%	11.7%	20.8%	25.2%	14.1%
Friday	7.1%	10.4%	16.9%	3.3%	4.1%	12.4%	17.5%	14.1%
Saturday	13.6%	9.3%	7.8%	20.9%	5.5%	7.7%	20.6%	14.1%

This table shows the probability of monitoring on a given day this week, conditional on the day monitored last week.

Figure W3: Distribution of monitored day of the week, by mode

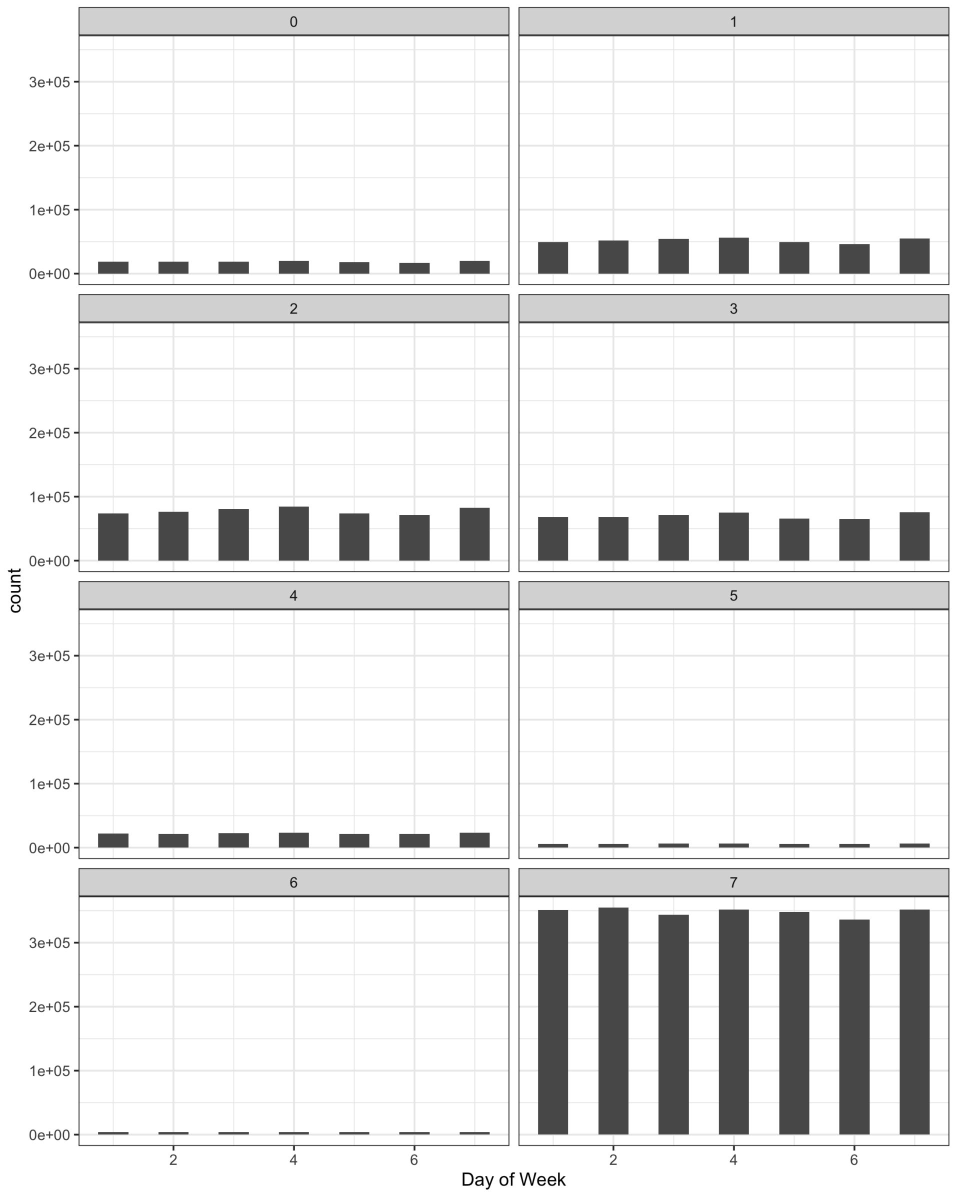

Figure W4: Distribution of monitored day of the month, by mode

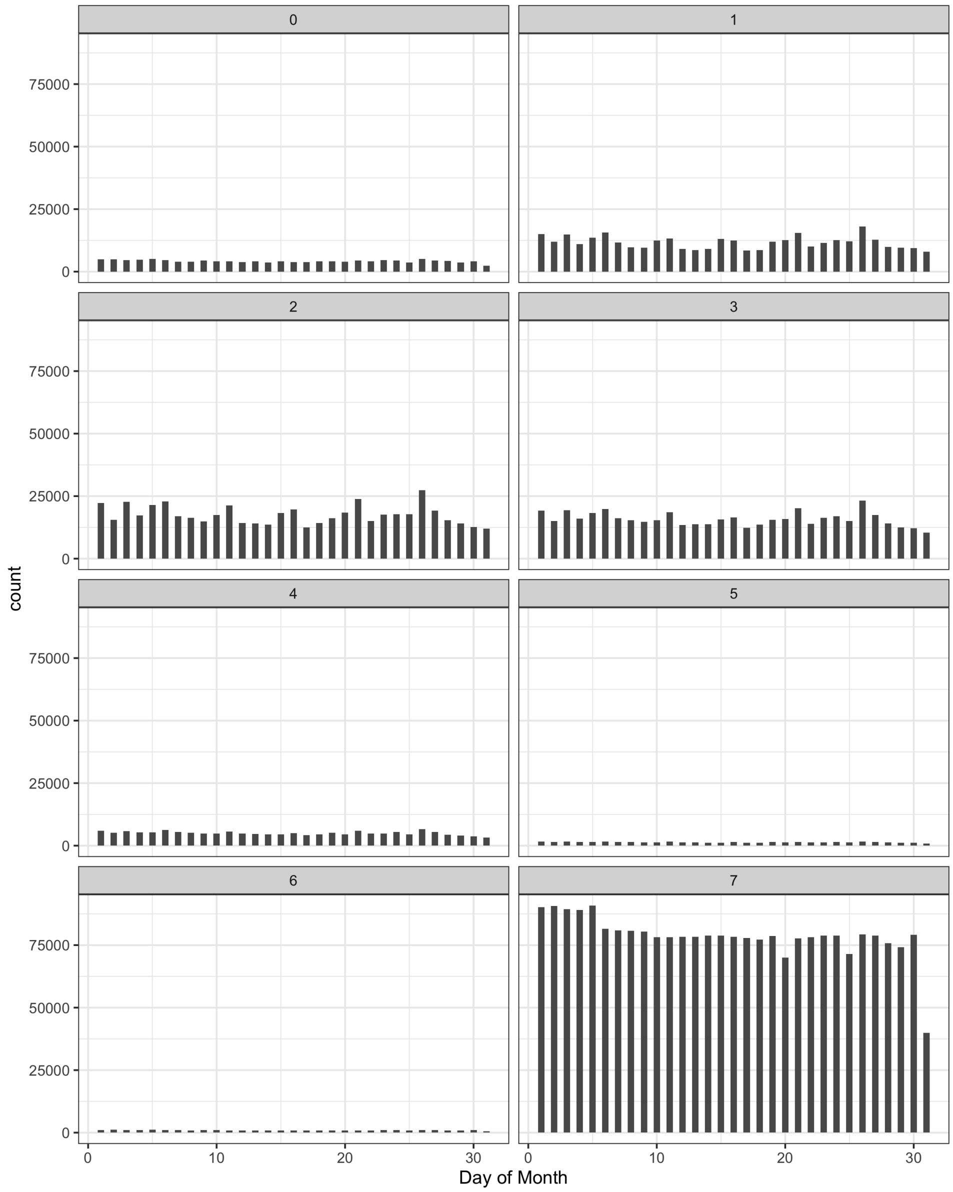

C Complementary Analysis for Section 4.4

C.1 Descriptive Evidence

To examine the focal retailer's price adjustment relative to price changes by competitor #1, we create a sample of the long-tail data that includes all the monitored prices of Competitor #1, which is a subset of *Group* 1 products. For each monitoring instance we log the price change by the competitor at that time, t, relative to the last observed price by the focal retailer, which can be at time $t-1$ or before. Then, we examine the focal retailer's price change relative to the price in day t, up to 7 days after the monitoring instance and log a "price response" if the focal retailer changed the price on day t within that period as well as the magnitude of the price change.[W2] To help isolate price adjustments associated with monitoring versus price adjustments associated with demand, we only include observations before the first sale of a product in our data (if any sales exist).

Table W3 presents the relationship between the competitor's price change and the focal retailer's price changes. In this table, we categorize price changes as: no change (the price is identical to the last observed price), a price reduction (price is lower than last observed price), or price increase (price is higher than last observed price) for the focal firm and Competitor #1. The results reveal several interesting phenomena. First, for these products in the long-tail for which the retailer monitors Competitor #1, the retailer engages in price exploration even when they observe no price changes by the competitor. The focal retailer changes the price within the week of monitoring 42.7% of the time after observing no competitor price change; the focal retailer is slightly more likely to increase prices (22.3%) relative to decreasing prices (20.4%). When the competitor is observed to change the price, the retailer is more likely to explore prices in the direction of the competing retailer's price change. Specifically, when a price decrease is observed, the proportion of price decreases increases by 5.6 percentage points (i.e., 26% versus 20.4%). When the competitor increases prices, the

[W2] We use a shorter time frame if the competitor was monitored again during the seven days and again changed the price, or if we reached the end of the observation period.

retailer is more likely to increase prices (28.4% versus 22.3%). In both cases, the adjustment of prices in the direction of the competitor's price change occurs roughly 27% more often compared to the steady state of no price changes by the competitor observed (i.e., 26/20.4 -1 = 0.27). This shows that the focal retailer adjusts price in the direction of the price change by Competitor #1. But, pricing behavior of the focal firm is still very stochastic.

Table W3: Price Change Matrix

	Focal Retailer (relative to day t)			
Competitor #1	Decrease	No change	Increase	Overall (%)
Decrease	29,458	59,168	24,705	113,331 (9.0%)
(row %)	(26.0%)	(52.2%)	(21.8%)	
No change	221,909	621,763	241,621	1,085,293 (85.8%)
(row %)	(20.4%)	(57.3%)	(22.3%)	
Increase	14,632	33,147	18,972	66,751 (5.3%)
(row %)	(21.9%)	(49.7%)	(28.4%)	
Overall (%)	265,999 (21.0%)	714,078 (56.4%)	285,293 (22.5%)	1,265,375

Note: The numbers without parentheses represent the number of observations in each cell. The numbers in parentheses represent proportion of observations. In the center of the matrix, they represent the fraction of observations in each row that correspond to that specific cell. For example, if Competitor #1 was observed to decreased the price in day t, then 26% of the time the focal retailer would also decrease the price, while they will increase the price 21.8% of the time. In the "Overall" row and column they represent the % of observations for each retailer that correspond to the three categorizations: price decrease, no change, and price increase. For example, the focal retailer doesn't change price 56.7% of the time, whereas Competitor #1 is observed to keep prices stable 85.9% of the time.

C.2 Event Study: Reaction to Monitoring

In this section, we construct an "event study" that allows us to isolate the impact of how the focal retailer adjusts price in response to competitive price information. A limitation of the previous analysis is that the focal retailer may be responding to other price changes by the competitor. To remove this confound, we consider long-tail products that have a price change by the largest competitor at day $t = 0$ after a prolonged period of no price changes. We only consider products that have no price changes by Competitor #1 the prior four weeks (28 days).[W3] Again, We remove the confound of past own sales by narrowing our analysis to Group 1 products that have zero sales from the beginning of the dataset through $t = 0$. This leads to 12,623 events and 10,004 different products.

These criteria lead to an event where the only likely explanation for a price adjustment is observing the competitor's price change at $t = 0$ (while the competitor may have changed prices before $t = 0$, they are only observable once monitoring occurred. As a reminder, in *Group* 1 monitoring is at a low-frequency cadence 66.8% of the time.). Since there are zero sales over the past four weeks from the retailer's own sales, the only possible feedback (about pricing) is that prices are too high. And, there is limited information from past competitor prices as they are constant. Because this setup ensures a pre-period of no changes by the competitor, it is reasonable to attribute a change in the focal retailer's behavior following the event to this recent change in the competitor's price.[W4] Thus, we treat the competitor's price change event as a plausible exogenous shifter and compare the focal retailer's behavior in the period before and after the event.

First, we examine the focal firm's frequency of price changes relative to t=0. In the pre-period of four weeks, the focal retailer changes prices 12.6% of the time. In Figure W1, we compare the rate of price changes in the 7 days after the event with the 7 days before the event. Day 0 represents the difference in frequency of price changes the day of the event compared with 7 days prior. We compute similar metrics for days 1 through 7 after the event. We observe a large increase of nearly 4 percentage points the day after the event; this

[W3] To construct the dataset, we ensure that the competitor was monitored by the focal retailer at least once in every one of the four weeks prior to the price changes, and that there is data for another four weeks after the price change.

[W4] It is still possible that both retailers change their prices at the same time due to another exogenous shifter, such as a change in cost, however, as we show most of the changes happen *after* Competitor #1 was observed to change prices and not before then.

difference persists for the next six days. We conclude from this that after the competitive price event there are more price adjustments (e.g., more price exploration) by the focal firm. In Table W4, we report a regression that shows that the average increase in price change frequency is 1.5% ($p < 0.01$) when controlling for product and day fixed effects. And, we find that there are more price changes when the largest competitor lowers their price (2.0%, $p < 0.01$) versus raises their price (1.0%, $p < 0.01$). These effects are also statistically different from each other ($p < 0.01$).

Figure W1: Change in frequency of price changes relative to the week before

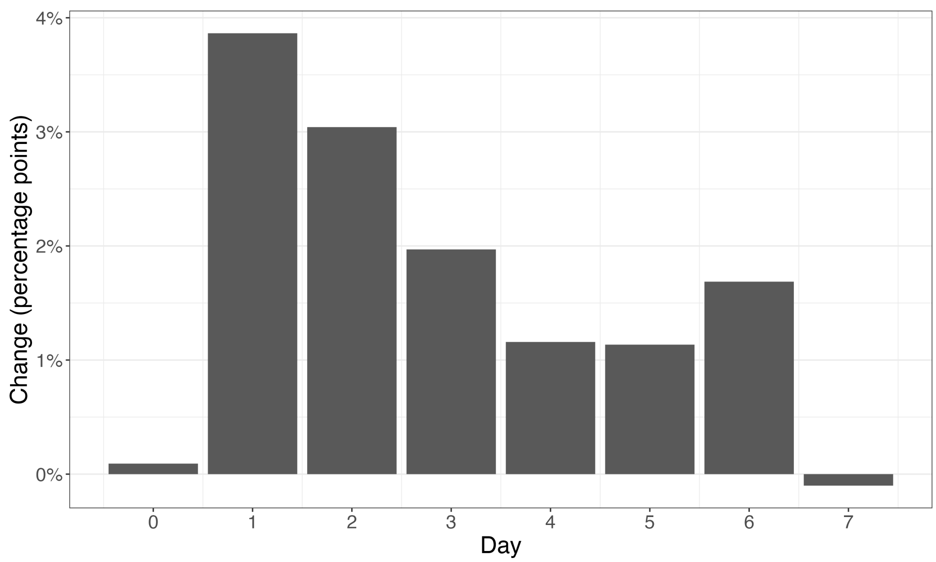

Next, we examine the magnitude of price changes after the event. If the largest competitor raises (lowers) its price on the event, we expect the focal firm to react in a similar direction. In Table W5, we show evidence consistent with this prediction by running regressions around the time of the event. Specifically, we use daily-product observations from 7 days before the event until the 7 days after the event of price change by the competitor, to estimate OLS models of the form:

$$Y_{it} = \alpha_{pos} After\_pos_{it} + \alpha_{neg} After\_neg_{it} + \beta X_{it} + \zeta_i + \eta_t \tag{1}$$

where Y_{it} is the log price of the focal retailer at time t for product i. To allow asymmetric response to price changes, we define $After\_pos$ as an interaction term to indicate that the

Table W4: Frequency of changes after an event

	(1)	(2)
After	0.015***	
	(0.002)	
After_pos		0.010***
		(0.002)
After_neg		0.020***
		(0.002)
Observations	187,456	187,456
Adjusted R2	0.311	0.311

Note: All regressions include product and day fixed effects. SE are clustered at the product level.

* significant at 10%; ** significant at 5%; *** significant at 1% level.

price change by Competitor #1 was a price increase; similarly, $After\_neg$ indicates a price decrease. X_{it} are control variables that include the price of Competitor #1 for product i in day t, and ζ_i and η_i are product and day fixed effects. To examine how the response depends on the magnitude of price change, we use the same model as Equation (1), but also add interaction terms of the absolute value of the rate of change by Competitor #1 to the $After$ variables.

In Column 1 of Table W5, we show that $After\_pos$ is positive and $After\_neg$ is negative (both are statistically significant), which suggests that the focal firm responds to the competitor's price as expected. In Column 2, we explore whether the change in the focal firm price is proportional to the price change by the largest competitor. Intuitively, one might expect that a larger competitive price change leads to a larger price adjustment. The coefficients for the interaction variables $After\_posXrate$ and $After\_negXrate$ confirm this intuition. Further, we see that the focal firm responds more to price decreases versus price increases. In Column 3, we add the log of the daily price of the competitor, log_price_C1, to the model. This allows us to test whether the focal firm is simply engaging in competitive price response (i.e., copying Competitor #1 price) or is exploring price adjustments. We find that the coefficient on the daily price of Competitor #1 is insignificant and the point estimate is very small. This suggests that the focal firm is not simply copying the competitor's price each day.[W5]

Together, the results of the event study provide convergent evidence with the descriptive analyses in Section C.1. When the largest competitor raises its price, the focal firm makes a price adjustment in the expected direction, but the response is very stochastic. The evidence also suggests the focal firm continues to engage in price exploration, but the focal firm explores lower (higher) prices after a price decrease (increase) by the competitor.

Table W5: Changes in Price After an Event

	(1)	(2)	(3)
After_pos	0.005***	-0.002	-0.002
	(0.001)	(0.002)	(0.002)
After_neg	-0.010***	-0.001	-0.001
	(0.001)	(0.002)	(0.002)
After_pos X rate		0.034***	0.039**
		(0.008)	(0.015)
After_neg X rate		-0.089***	-0.097***
		(0.015)	(0.022)
log_price_C1			-0.009
			(0.015)
Observations	187,456	187,456	187,456
Adjusted R2	0.988	0.988	0.988

Note: All regressions include product and day fixed effects. SE are clustered at the product level.
* significant at 10%; ** significant at 5%; *** significant at 1% level.

[W5] Note that the competitor's price is constant for 7 days prior to the event but may change after the event.

D Complementary Analysis for Section 5

In this section we include four additional analyses: (i) replicating the of Table 5 using the sales index instead of the revenue index as the dependent variable; (ii) using a shorter post-period that excludes the holidays (post period of two month, ending in October 31, 2019); (iii) using daily level data instead of two time periods; and (iv) exploring the association between past sales of similar products and future revenues.

D.1 Using sales instead of revenues

Table W6 repeats the analysis in Table 5, using the sales index as the outcome variable.

Table W6: Monitoring and Sales

	(1)	(2)	(3)	(4)	(5)	(6)
Post	2.214***	3.070***	8.718***	3.221***	4.143***	7.644***
	(0.591)	(0.560)	(0.897)	(0.517)	(0.494)	(0.672)
Post X Monitor	9.260***	10.129***	4.278***			
	(0.734)	(1.237)	(1.449)			
Post X Med Monitor				7.552***	7.303***	3.811
				(2.837)	(2.678)	(2.723)
Post X High Monitor				67.914***	67.774***	64.299***
				(1.820)	(8.994)	(9.081)
Mean Price	0.0001	-0.269***	-0.209***	-0.00004	-0.210***	-0.180***
	(0.001)	(0.061)	(0.075)	(0.001)	(0.051)	(0.065)
Price Change Frequency	15.877***	12.533	13.447	14.234***	10.489	18.085
	(4.636)	(16.029)	(12.073)	(4.599)	(14.865)	(11.659)
Price Variation	12.927**	-9.142	-9.025	17.832***	5.346	-1.334
	(5.290)	(18.090)	(15.639)	(5.211)	(17.824)	(15.391)
# Obs	92,428	92,428	92,428	92,428	92,428	92,428
Adjusted R2	0.004	0.005	0.005	0.017	0.017	0.017
Product FE	No	Yes	Yes	No	Yes	Yes
Matching	No	No	Yes	No	No	Yes

[*] significant at 10%; ** significant at 5%; *** significant at 1% level.

D.2 Replication using daily-level observations

We replicate our results also using daily level data for the same products as the main analyses. We use regressions that interact pre-regression characteristics (monitoring, high frequency of changes, monitoring Competitor #1, number and sales of similar products), and product and time FE:

$$Y_{it} = \alpha Post_t \times Monitor_i + \beta Post_t \times HighFreq_i + \gamma Post_t \times Monitor_i \times HighFreq_i + \eta X_{it} + \zeta_i + \eta_t \quad (2)$$

In these specification, X_{it} includes the price and monitoring of product i at time t, and the price change rate from day $t-1$ to day t.

Table W7 presents the main results, with the unmatched and matched samples. The results are consistent with the results in section 5.2.

D.3 Replication using a shorter post-period

Next, we replicate the main analyses using a shorter time period, ending at the end of October 2019. We replicate all three tables from section 5, which are Tables W8, W9, W10. Overall, while the coefficients are smaller in magnitude, our main results and conclusions hold. The only result that is not consistently replicated is the result of monitoring Competitor #1. In Table W10, the coefficient of the interaction with the Competitor #1 monitoring is not statistically significant. This is likely because a substantial amount of sales happened during the holiday period.

D.4 Past Sales of Similar Products and Future Sales

Since demand signals are rare in the long-tail, another source of information is sales of similar products. For example, a product may be sold in multiple colors, styles and sizes and it may be possible to learn about demand for a focal product by looking at similar products. This is consistent with an algorithm proposed by Mussi et al. (2022) to learn about demand for long-tail products. Their idea is that sales of similar items can inform a retailer about

Table W7: Daily-level results

	(1)	(2)	(3)	(4)	(5)	(6)
Post X Monitor	0.782***	0.668***	0.627***	0.514***	0.707***	0.697***
	(0.133)	(0.133)	(0.126)	(0.127)	(0.119)	(0.135)
Post X High Frequency	0.109***	0.215***	0.109***	0.215***	0.102***	0.210***
	(0.028)	(0.070)	(0.028)	(0.070)	(0.029)	(0.072)
Post X Monitor X High Freq	-0.594***	-0.699***	-0.609***	-0.713***	-0.487***	-0.613***
	(0.136)	(0.158)	(0.138)	(0.160)	(0.123)	(0.149)
Post X Monitor Comp #1			0.200**	0.199**		
			(0.089)	(0.088)		
Post X Sales Similar Products					0.0001***	0.0001***
					(0.00003)	(0.00002)
$Price_t$	-0.021***	-0.014***	-0.021***	-0.014***	-0.020***	-0.013***
	(0.006)	(0.005)	(0.006)	(0.005)	(0.006)	(0.005)
$Monitor_t$	0.043	0.060**	0.044	0.061**	0.044	0.060**
	(0.027)	(0.028)	(0.027)	(0.028)	(0.027)	(0.028)
Price Change Rate	0.108***	0.072**	0.108***	0.072**	0.101***	0.067**
	(0.036)	(0.028)	(0.036)	(0.028)	(0.034)	(0.026)
# Obs	8,769,643	8,769,643	8,769,643	8,769,643	8,769,643	8,769,643
Adjusted R2	0.023	0.022	0.023	0.022	0.023	0.022
Matching	No	Yes	No	Yes	No	Yes

Note: All regressions include product and week fixed effects. SE are clustered at the product and day level.
* significant at 10%; ** significant at 5%; *** significant at 1% level.

demand for low demand items.

To capture this in our analysis we add another interaction: the total (indexed) number of sales of the products from the same brand and sub-category in the pre-period. There is substantial variation in the (indexed) sales of similar products: while 21.5% of the products in this subsample have no sales of similar product, the median indexed sales for similar products is 656 and the standard deviation is 12,838.

Table W11 presents the results. Column (4) suggests that relative to a product that does not have similar product sales, a median sales of similar products (656) is associated with 4.592 additional indexed revenue units. This suggests that similar products that are

Table W8: Monitoring and Demand

	(1)	(2)	(3)	(4)	(5)	(6)
Post	3.201***	3.671***	15.331***	4.943***	5.481***	12.414***
	(0.755)	(0.652)	(3.971)	(0.666)	(0.598)	(2.355)
Post X Monitor	11.045***	11.376***	-0.748			
	(0.928)	(1.553)	(4.502)			
Post X Med Monitor				34.398***	34.255***	27.130**
				(3.608)	(13.252)	(13.424)
Post X High Monitor				58.728***	57.692***	50.241***
				(2.353)	(10.257)	(10.540)
Mean Price	0.008***	-0.241**	-0.060	0.008***	-0.185**	-0.039
	(0.002)	(0.098)	(0.304)	(0.002)	(0.094)	(0.301)
Price Change Frequency	3.976	17.665	23.144	0.562	8.884	27.623
	(5.444)	(10.798)	(41.849)	(5.421)	(10.125)	(40.977)
Price Variation	-29.617***	-40.942***	-54.753***	-19.948***	-25.649***	-49.900***
	(6.520)	(10.707)	(17.885)	(6.471)	(7.632)	(18.601)
# Obs	92,214	92,214	92,214	92,214	92,214	92,214
Adjusted R2	0.003	0.004	0.003	0.009	0.009	0.005
Product FE	No	Yes	Yes	No	Yes	Yes
Matching	No	No	Yes	No	No	Yes

* significant at 10%; ** significant at 5%; *** significant at 1% level.

more popular could serve as a useful signal for long-tail products. At the same time, we observe that the monitoring and frequent price changes coefficients are largely unchanged. The stability of these coefficients suggests that demand from similar products may be an independent signal of future demand. But, the magnitudes of the coefficients suggest that there are much stronger associations with monitoring and past price changes.

Table W9: Monitoring and Price Changes

	(1)	(2)	(3)	(4)
Post	1.458	2.320***	8.263***	7.382***
	(0.921)	(0.606)	(1.562)	(2.151)
Post X Monitor	20.261***	20.155***	13.407***	13.457***
	(1.488)	(3.154)	(2.373)	(4.541)
Post X High Frequency	5.327***	4.657***	13.390***	14.190**
	(1.267)	(1.417)	(1.827)	(6.898)
Post X Monitor X High Freq	-16.064***	-15.143***	-24.010***	-23.400***
	(2.261)	(3.768)	(3.352)	(8.700)
Mean Price	0.008***	-0.238**	0.011***	-0.050
	(0.002)	(0.098)	(0.002)	(0.308)
Price Change Frequency	3.639	10.222	9.959	32.937
	(5.889)	(10.964)	(7.033)	(39.434)
Price Variation	-28.320***	-37.106***	-32.392***	-51.619***
	(6.521)	(10.133)	(8.132)	(17.100)
Monitor X High Frequency	0.199		0.087	
	(1.067)		(1.551)	
# Obs	92,214	92,214	92,214	92,214
Adjusted R2	0.004	0.004	0.004	0.004
Product FE	No	Yes	No	Yes
Matching	No	No	Yes	Yes

* significant at 10%; ** significant at 5%; *** significant at 1% level.

Table W10: Monitoring Competitor #1

	(1)	(2)	(3)	(4)
Post	1.458	2.312***	8.263***	7.373***
	(0.921)	(0.604)	(1.562)	(2.150)
Post X Monitor	18.162***	17.454***	11.339***	10.575**
	(2.079)	(3.152)	(3.243)	(4.650)
Post X Monitor Comp #1	2.711	3.486	2.670	3.724
	(1.875)	(3.257)	(2.855)	(3.263)
Post X High Frequency	5.327***	4.660***	13.390***	14.196**
	(1.267)	(1.417)	(1.827)	(6.898)
Post X Monitor X High Freq	-16.274***	-15.391***	-24.216***	-23.678***
	(2.266)	(3.856)	(3.359)	(8.734)
Mean Price	0.008***	-0.238**	0.011***	-0.050
	(0.002)	(0.098)	(0.002)	(0.308)
Price Change Frequency	3.642	10.644	9.959	33.206
	(5.889)	(10.931)	(7.033)	(39.445)
Price Variation	-28.343***	-37.774***	-32.408***	-52.092***
	(6.521)	(10.232)	(8.132)	(17.159)
Monitor X High Freq	0.199		0.087	
	(1.067)		(1.551)	
# Obs	92,214	92,214	92,214	92,214
Adjusted R2	0.004	0.004	0.004	0.004
Product FE	No	Yes	No	Yes
Matching	No	No	Yes	Yes

* significant at 10%; ** significant at 5%; *** significant at 1% level.

Table W11: Sales of Similar Products

	(1)	(2)	(3)	(4)
Post	-9.050***	-3.073	-5.123	-3.874
	(2.945)	(4.755)	(3.459)	(6.244)
Post X Sales Similar Products	0.008***	0.008***	0.007***	0.007***
	(0.0002)	(0.002)	(0.0002)	(0.002)
Post X Monitor	72.820***	73.567***	70.784***	71.475***
	(4.756)	(10.003)	(5.230)	(11.051)
Post X High Frequency	11.358***	8.173*	20.237***	21.689***
	(4.043)	(4.255)	(4.039)	(6.900)
Post X Monitor X High Freq	-51.204***	-51.071***	-60.661***	- -60.960***
	(7.244)	(11.743)	(7.423)	(12.949)
Mean Price	0.025***	-1.125***	0.022***	-0.824**
	(0.005)	(0.393)	(0.005)	(0.354)
Price Change Frequency	1.767	-33.053	21.692	36.229
	(20.501)	(64.450)	(16.839)	(46.935)
Price Variation	10.095	-10.658	3.672	-0.434
	(21.489)	(40.342)	(19.477)	(38.699)
Monitor X High Frequency	0.094		-0.446	
	(3.455)		(3.448)	
# Obs	92,428	92,428	92,428	92,428
Adjusted R2	0.028	0.029	0.027	0.028
Product FE	No	Yes	No	Yes
Matching	No	No	Yes	Yes

* significant at 10%; ** significant at 5%; *** significant at 1% level.

E Matching Algorithm

To match between the two long-tail groups, we use generalized full matching (Sävje et al. 2021). We implement it using the quickmatch package in R (Sävje et al. 2018). We use the following pre-period variables for matching between products: mean price, price change frequency, price variation coefficient, MAP policy, revenue from the brand, revenue from the brand and sub category of the product, and the number of products with the same brand and sub category for this product. When matching, we compute the propensity of the product to have any monitoring (i.e., belong to *Group* 1).

Figure W5 presents the standardized mean difference between the variables we used for matching. Importantly, the distance between both groups is indistinguishable from zero, and most of the variables lie within the default intervals for acceptable balances (indicating differences of .05 and .1).

Figure W5: Matching: Standardized mean differences

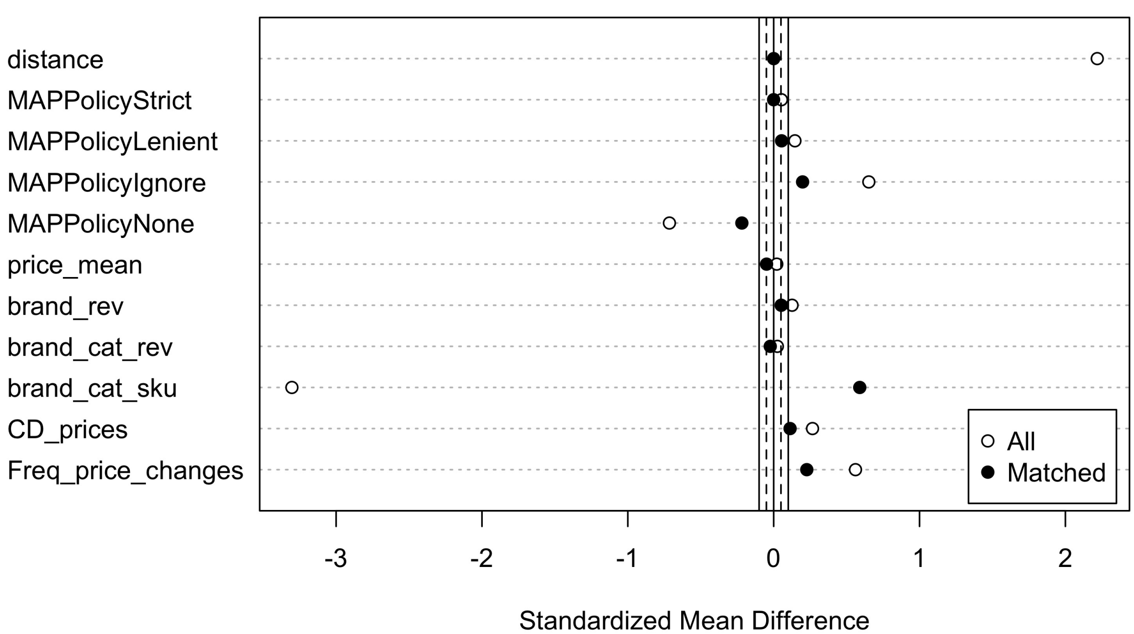

Figure W6a presents the histograms of the propensity scores before and after the match, and Figure W6b presents the relative weights implied by matching for each product. We then use those weights as an input to the regressions we run using the matched sample.

Figure W6: Matching: Propensity scores

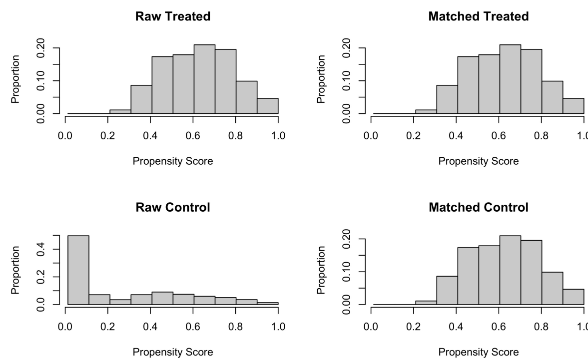

(a) Histograms

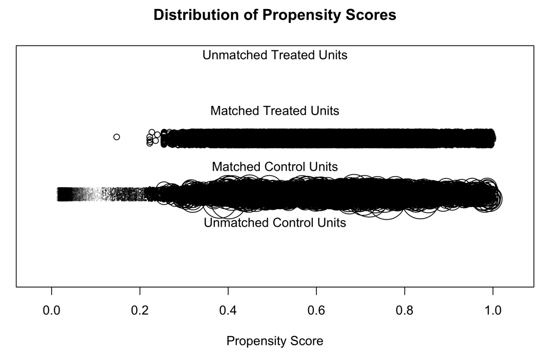

(b) Implied regression weights

Finally, Figure W7 includes density plots for each of the variables to allow visual assessment of the match.

Figure W7: Matching: Density plots