Competition-Based Dynamic Pricing in Online Retailing

Marshall Fisher
The Wharton School, University of Pennsylvania, fisher@wharton.upenn.edu

Santiago Gallino
Tuck School of Business, Dartmouth College, santiago.gallino@tuck.dartmouth.edu

Jun Li
Ross School of Business, University of Michigan, junwli@umich.edu

Jerry Liu
Head of Pricing and Category Management, Yihaodian, jerry.liu@yhd.com

Gang Yu
Co-Founder and CEO, Yihaodian, gang@yhd.com

1. Executive Summary

We partner with Yihaodian[1], a leading Chinese online retailer, to tackle the challenge of real-time competitive pricing in online retailing. Through this year-long close partnership, we

- Design and estimate an empirical model to capture consumer choices among substitutable products from multiple retailers. This model provides a key parameter: the degree to which consumers engage in price comparison across competing retailers, a key factor in determining the optimal level of competitive response.
- Deploy a novel identification strategy to address the methodological challenge that almost all retailers face in estimating a consumer choice model: the absence of competitor sales information. This strategy exploits Yihaodian's and their competitors' stock-outs as a valid source of variation to the consumer choice set.
- Successfully introduce price variations through a randomized fractional factorial design, which lays the foundation for accurate and unbiased estimates of product-specific price elasticities.
- Achieve an average MAD of 37.7% in predicting *daily* SKU sales, and even lower MAD (26.1%) for high volume products. We also find that incorporating competitor stock-outs significantly improves demand estimation accuracy.
- Implement a best-response pricing algorithm that takes into account consumer choice behavior, competitors' actions, and supply parameters (procurement costs, margin target, and manufacturer price restrictions) at the partner retailer. Through a carefully controlled filed experiment in baby feeding bottles, we document an *eleven percent revenue increase*, while maintaining margin above a specified target set by the retailer.

The algorithm has then been integrated with Yihaodian's IT system, and they are implementing a phased roll-out to other categories. Meanwhile, the existing price tracking system has been re-evaluated and configured to ensure accurate real-time data gathering.

2. The Challenge of Competitive Pricing in Online Retailing

The Internet has changed the way price information is disseminated. With just a few clicks consumers are able to obtain price information from multiple retailers. This increased price transparency induces fierce competition among online retailers and requires real-time monitoring and quick responses to competition.[2]

The price transparency enjoyed by consumers has prompted many online retailers to adopt a competition-based pricing strategy in which they constantly monitor competitors prices and use this as an input in setting their own prices. For example, they may always charge x dollars or

[1] http://www.yhd.com

[2] Coming soon: Toilet paper priced like airline tickets. The Wall Street Journal. September 5, 2012.

x percent lower or higher than a target competitor or any competitor with the lowest price. Not surprisingly, retailers miss several opportunities with such simple heuristics. Instead, they should ask themselves, shouldn't my reaction depend on consumers' elasticity to prices? Shouldn't my reaction depend on the extent to which consumers compare prices across retailers and stay loyal to a retailer (e.g. postpone purchase or substitute to a similar product from the same retailer)? Shouldn't my reaction depend on changes in availability at competing retailers? Should I still match prices if it seems like the competitor made a pricing mistake? We address exactly these questions in this paper.

Determining the best-response price requires knowing how demand reacts to price changes. This is a challenging task. Simply regressing historical sales on prices while controlling for observable product characteristics and seasonality usually suffers from endogeneity issues. Pricing managers often observe demand signals that we researchers do not, such as unobserved product characteristics or a temporal surge in demand due to manufacturer advertisements, and they may adjust prices based on observed demand signals. If they increase price when they see a demand surge, this creates a correlation that fallaciously implies a higher price results in higher demand. Moreover, the relationship between demand and price will be further confounded by the price levels of substitutable products that the same retailer offers and the price levels of the same product that the competition offers.

To determine the best-response price we also need to understand the extent to which consumers compare prices across retailers. In the situation where consumers are perfectly loyal to their choice of retailers—that is, they will only substitute within a retailer but not across retailers—there is no need to match competitor price changes to any extent. However, in the situation where consumers always choose the cheapest retailer for any product they buy, we need to either charge the lowest price in the market or accept no sales. Accurately assessing the level of consumer engagement in price comparison across retailers will allow targeted price responses that are efficient and effective.

We partnered with a leading Chinese online retailer—Yihaodian, which we will refer to as the retail partner hereafter—to address these challenges. Our retail partner, Yihaodian, is a leading Chinese online retailer that originally focused on consumer packaged goods but over time evolved to be a hypermarket. Yihaodian was founded in July 2008 and achieved sales of $1.9 billion in 2013. A 2011 survey by Deloitte[3] identified Yihaodian as the fastest growing technology company in the Asia-Pacific region, with a three year revenue growth of 19,218 percent.

[3] Deloitte News Release: Top 10 Fastest Growing Technology Firms for 2011. December 1, 2011.

3. Our Solutions

Our solution approach can be divided into four stages that utilize different methodologies, including structural modeling and estimation, experimentation, and optimization. These stages are closely connected in the sense that each stage provides necessary inputs to inform the next one.

First, we developed a demand model to understand how consumers make choices when given a set of substitutable products from multiple retailers, the estimation of which can provide inputs to determine optimal responses to competitors' price and availability variations. Our model resolves a key challenge many retailers face when attempting to implement a choice model to understand consumer purchase decisions: absence of competitor sales information. We propose a novel identification strategy that exploits temporary variations in consumer choice sets through our own and competitor stock-outs. These variations provide the additional moment conditions necessary to estimate consumer preferences of retailers and their level of engagement in comparing prices across competing retailers.

Second, we conduct a randomized price experiment to obtain unbiased estimates of price elasticities. We randomly assign prices to each product under study using a fractional factorial design. We focus on one particular product category sold by this retailer: baby-feeding bottles. This category presents a number of features that make it very attractive for our study. First, it includes a group of relatively homogeneous products that can be characterized by a small number of well-defined product attributes. Second, the life cycle of the products is long compared to the time span that the product will be used. Third, the associated brands do not engage in exclusivity deals with retailers. Finally, most customers do not engage in repeated purchases in a short period of time (e.g., daily or weekly) since the product will outlast the baby's need. Therefore, inter-temporal substitution is not a pressing concern. It is also worthwhile to note these features are not unique only to baby feeding bottles. There are many other product categories that share similar characteristics and where our methodology and analysis also apply, such as small appliances, hardware tools, and kitchenware, to name a few.

It is also important to note that with the presence of competition, price randomization alone will not necessarily guarantee unbiased estimation of elasticities unless competitors' actions are properly accounted for. Ignoring competitors' reactions to our price changes would bias the estimation because prices can still be correlated with unobserved demand shocks through correlation with competitors' prices. This is why we account for changes in competitors' prices and product availability in the consumer choice model.

Third, we solve a constrained optimization problem to define optimal price responses to competitor price changes. Based on the estimates obtained from the choice model and taking into account competitor prices and product availability, we find prices for our partner retailer that maximize total revenue for the category subject to several constraints imposed by the retailer. The constraints include a lower bound on average category margin, lower and upper bounds on individual product margins, and manufacturer price restrictions.[4] Changes in recommended prices come from four different sources: (1) changes in costs, (2) changes in our own product availability, (3) changes in competitor prices, and (4) changes in competitor product availability.

Last, in collaboration with our retail partner, we evaluate the performance of our best-response pricing algorithm through a carefully controlled field experiment. In order to evaluate the impact on total category revenue, instead of matching products based on product features we assign products to treatment and control groups to minimize substitution across groups but meanwhile allow substitution within each group. We apply our algorithm in one geographical region where the retailer operates and choose two other similar but disparate regions where the retailer also operates as comparison. This design leads to a Difference-in-Differences-in-Differences estimator, which allows us to correct for the potential differences in demand trends between the control and treatment groups with the presence of comparison groups subject to similar but independent demand.

4. Results and Impact

Results. We compare estimation results obtained from different types of data, observational vs. experimental, and different types of models, with and without competition and price comparison, in Table 1. While the attached paper provides detailed explanations of these results, we would like to draw your attention to the parameter λ ($0 \leq \lambda < 1$) in Column 4, which captures the extent to which consumers conduct price comparison. The larger the value of λ, the more intense the price comparison is across competing retailers. The estimate of λ equals 0.7911, which suggests consumers engage in extensive price comparisons across competitors. The high intensity of price comparison indicates that retailers need to follow competitors' price adjustments closely to stay competitive in the market. The impact of either overpricing or underpricing can be significant.

We next demonstrate the importance of considering the degree to which consumers compare prices across retailers. Substitution patterns are the key to responsive pricing. On one extreme, if customers do not substitute across retailers, there is no need to follow competitors' prices; on the other extreme, if customers always compare prices, one should almost always follow competitors' price changes. From our model's estimation, consumers exhibit a strong price comparison behavior. Therefore, prices should be, and indeed are, more responsive to competitor price changes under such condition (compare the differences between Column (5) and (6) to differences between Column (3) and (4) in Table 2). A model that fails to capture the extent to which consumers compare prices

[4] We omit details of these constraints for confidentiality reasons.

across retailers will lead to insufficient response to competitor price changes. This will be true even when we are able to capture price elasticities without bias with a randomized price experiment.

Figure 1 shows the goodness of fit graphically. It plots the predicted daily sales (in green) against observed daily sales (in blue). The average daily Mean Absolute Deviation (MAD) is 0.377. Note there is a negative correlation of -0.697 (sig= 0.0041) between model goodness of fit by product, as measured by MAD, with average daily sales. In other words, the model better predicts demand for fast-moving products than for slow-moving products.

Impact. Implementation of our solution led to *11 percent increase* in total category revenue. We demonstrate this through a carefully controlled live experiment.

The objective of our pricing algorithm is to maximize revenue for the category. Price changes of a specific product will lead to not only revenue changes for that particular product but also other products with similar features due to substitution. For this reason, a valid experiment design requires minimal substitution between treatment and control groups, otherwise the control group will be contaminated due to the spillover effect. Instead of matching products on their main attributes, we identify that one existing attribute that allows a clean separation of the market segments: bottles designed for certain ranges of babies' ages.

Our retail partner operates in multiple regions of the country. These regions are geographically separated, and are supported by separate distribution and pricing teams. We implement the pricing algorithm in only one of the regions, using two other regions as controls (chosen based on similarity in population, income level and retailer targeted margins). The experiment design we implemented allows us to adopt a triple-differences estimator to measure the impact of the proposed pricing methodology. The triple differences come from comparisons of the periods before and after, different regions, and the treatment and control age groups. Table 3 shows that the revenue increases for the treated category vary from 10.9 percent to 12.4 percent depending on control variables included in the regression.

This revenue improvement is not unique to baby feeding bottles. We are currently expanding the implementation of the algorithm to kitchenware and small appliances. Based on our preliminary analysis of kettles, we obtained 19 percent revenue improvement in this category.

5. Conclusion

Our work contributes to the practice of revenue management and pricing in a number of ways. First, we propose a parsimonious choice model that captures the key tension involved in this competitive environment, and we propose a novel identification strategy using own and competitor stock-outs to provide additional moment conditions in the absence of competitor sales information.

Second, we conduct a randomized price experiment in the field to obtain unbiased measures of price elasticities, thereby overcoming the limitations of using observational data. We provide examples to illustrate several problems with observational data. Estimates of price elasticities can show up as statistically insignificant from zero, that is, non-distinguishable from inelastic demand, due to lack of price variations historically, which happens very often when selling millions of products online. Sometimes, even if the price of a product itself varied historically, it follows closely competitors prices such that there is little price variation comparatively. In this situation, it is impossible to distinguish how demand responds to changes in one retailer's own price versus changes in the competition price. Moreover, estimates of price elasticities can be biased upward when ignoring the fact that retail managers make price decisions based on private demand signals. Duncan Painter, the CEO of WGSN Group, a firm specializes in fashion forecasting for retailers, commented that they often use price as a proxy for sales—discounted price implies low sales and vice versa, precisely due to this reason[5]. As Uri Gneezy and John List pointed out, "running [field] experiments is a costly undertaking, but it is prohibitively costly *not* to experiment.". In fact, "many product and pricing failures can be laid at the feet of insufficient investigations and tests."

Third, our methodology has stood the test of a real competitive business environment and demonstrated tangible revenue improvements. Working closely with the industry partner to test our methodology in the field, we are able to learn not only whether the proposed methodology improves business decisions, but also, perhaps more importantly, the challenges and opportunities that implementing a competition-based dynamic pricing policy can bring to an online retailer in a real setting. Our work helps navigate and evaluate the trade-offs involved in bridging theoretical, empirical, and field work.

Finally, our work presents a scalable and replicable methodology to set dynamic prices in an online retail setting. In what follows, we present a detailed account of our methodology and the results.

[5] From a private communication with Duncan Painter.

Table 1 Estimates of Product Elasticities for Feeding Bottle SKUs

	Historical Data	Randomized Price Experiment		
	(1)	(2)	(3)	(4)
	Multinomial Logit w/o Competitor Info	Multinomial Logit w/o Competitor Info	Full Choice Model w/ Competitor Info Indep. Demand Shocks	Full Choice Model w/ Competitor Info Corr. Demand Shocks
Product Specific Price Coefficients ($\beta_j^{\dagger}$)				
Product 1	0.9742*	-5.0330***	-5.4379***	-1.6747***
Product 2	-2.4509***	-1.5239*	-2.2669**	-0.3667***
Product 3	-5.4356***	-7.1360***	-4.5718***	-6.7734***
Product 4	-4.3981**	1.9038*	-0.0056	-0.0036
Product 5	-1.8482**	-0.1298*	-0.0039	-0.9532
Product 6	-1.4099*	-3.6913*	-2.3672***	-1.0537***
Product 7	-0.2136*	-1.6313*	-2.3989***	-0.5404***
Product 8	3.7093*	-4.3830***	-2.9645***	-1.1644***
Product 9	-3.7126***	-4.3691***	-4.2956***	-1.1176***
Product 10	2.2436*	-4.3312*	-4.8453***	-4.1492***
Product 11	-8.2248***	-4.0908***	-4.0576***	-0.5038***
Product 12	1.4870**	-5.2959*	-3.0487***	-2.1872***
Product 13	0.8118*	-19.9489***	-4.5145***	-11.2814***
Product 14	-3.6317***	-2.1492**	-2.5254***	-0.9216***
Product 15	-2.9135***	-3.7712***	-3.8478***	-1.1421***
Retailer Preferences (α_r)				
Retail Partner	-	-	0 (baseline)	0 (baseline)
Competitor 1	-	-	2.5984	0.2172
Competitor 2	-	-	0.3658	0.0169
Competitor 3	-	-	-3.7499	-1.8363***
Competitor 4	-	-	-3.8870***	-2.4642**
Extent of Price Comparison (λ)	-	-	-	0.7911***
Product Specific Intercepts (α_j)	Yes	Yes	Yes	Yes
Day of Week, Holiday Dummies (γ)	Yes	Yes	Yes	Yes
# Days	67	30	30	30
# Purchases	7690	3742	3742	3742
Pseudo R-Square	0.6696	0.5848	0.7107	0.7164
Mean Absolute Deviation (MAD)	0.419	0.480	0.382	0.377
Log Likelihood	-35973.844	-17041.335	-17053.73	-17036.01

†: In Columns 1 - 3, $\text{elasticity}_{jA} = \beta_j(1 - s_{jA})$, and $\approx \beta_j$ if s_{jA} is small. In Column 4, $\text{elasticity}_{jA} = \frac{\beta_j}{1-\lambda}(1 - s_{jA}) - \frac{\lambda}{1-\lambda}\beta_j(s_{jA} + ns_{jA})$, where ns_{jA} denotes the share of product j offered by retailer A within the product nest, i.e., the set of product j's available at all retailers.

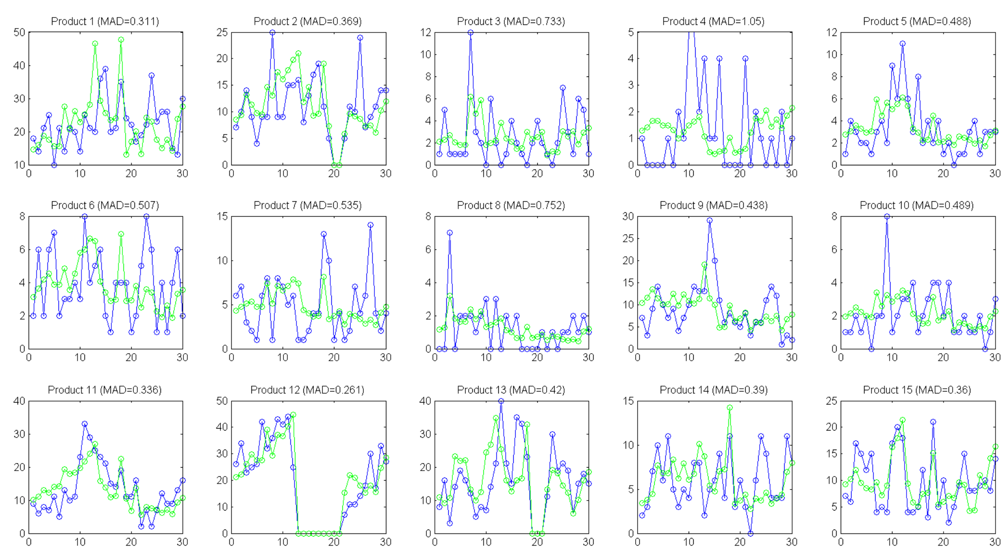

Figure 1 Model Goodness of Fit

Table 2 Best-Response Prices under Different Models

	Competitor 1 Price		Suggested Own Price w/o Price Comparison		Suggested Own Price w/ Price Comparison	
	(1)	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**
Product #	**Day 1**	**Day 2**	**Day 1**	**Day 2**	**Day 1**	**Day 2**
1	42.8	42.8	38.1	37.4	39.1	39.2
2	36.8	36.8	38.6	38.4	37.5	36.8
3	48.8	48.8	46.8	47.8	45.0	45.0
4	48.8	48.8	49.8	48.3	48.3	48.5
5	108.0	108.0	87.7	87.8	87.5	86.3
6	108.0	108.0	87.4	87.3	89.3	87.9
7	79.0	79.0	81.7	81.5	82.6	82.6
8	85.0	85.0	72.3	72.3	74.4	74.2
9	108.0	108.0	84.4	84.4	87.9	86.6
10	18.0	18.0	17.0	16.7	18.9	18.9
11	108.0	108.0	100.6	100.6	101.5	101.5
12	98.0	98.0	79.1	80.3	81.0	81.4
13	98.0	98.0	80.9	81.3	83.2	83.4
14†	**105.0**	**95.0**	**98.8**	**97.8**	**103.1**	**97.5**
15†	**109.0**	**119.0**	**104.7**	**105.9**	**99.3**	**103.2**

†: Competitor 1 changes prices for product 14 and 15 from day 1 to day 2 while keeping other prices unchanged.

Table 3 Revenue Impact of Best Response Pricing

ln(daily revenue)	(1) w/o daily margin and traffic	(2) w/ daily margin	(3) w/ daily margin & traffic
Treatment (α_3)	**0.109***	**0.112***	**0.124***
Group 1 dummy	-0.118***	-0.114***	-0.117***
No test week 1	Baseline		
test week 1	-0.284***	-0.283***	-0.283***
test week 2	-0.112	-0.114	-0.124*
no test week 2	0.267***	0.270***	0.246***
test week 3	0.138**	0.136**	0.100
Location 2	-0.656***	-0.779***	-0.977*
Location 3	-1.257***	-1.282***	-2.026***
daily margin		-0.539	-0.363
Location2 X daily margin		1.683	1.504
Location3 X daily margin		0.314	0.402
ln(daily traffic)			0.020
Location2 X ln(daily traffic)			0.062
Location3 X ln(daily traffic)			0.216
day of week dummy	yes	yes	yes
month dummy	yes	yes	yes
const	7.976***	8.017***	7.922***
# obs	432	432	432
# treatment	38	38	38
R-sq	0.7322	0.733	0.7369

Note: Products and dates are dropped when the algorithm was not correctly implemented.