How Do Price Promotions Affect Customer Behavior on Retailing Platforms? Evidence from a Large Randomized Experiment on Alibaba

Dennis J. Zhang[1]*, Hengchen Dai[2]*, Lingxiu Dong[1],
Fangfang Qi[3], Nannan Zhang[3], Xiaofei Liu[3], Zhongyi Liu[3], Jiang Yang[3]
1. Olin Business School, Washington University in St. Louis
2. Anderson School of Management, University of California, Los Angeles
3. Alibaba Group Inc.

Dynamic pricing through price promotions has been widely employed by online retailers. We study how a promotion strategy–offering customers a discount for products in their shopping cart–affects customer behavior in the short and long term on a retailing platform. We conducted a randomized field experiment involving more than 100 million customers and $11,000$ retailers with Alibaba Group, the world's largest retailing platform. We randomly assigned eligible customers to either receive promotions for products in their shopping cart (treatment group) or not (control group). In the short term, our promotion program doubled the sales of promoted products. In the long term, we causally document unintended consequences of this promotion program during the month following our treatment period. On the positive side, it boosted customer engagement, increasing the daily number of products customers viewed and their purchase incidence on the platform. On the negative side, it intensified strategic customer behaviors in the post-treatment period in two ways, by (1) increasing the proportion of products that customers added to the shopping cart upon viewing them, possibly due to their anticipation of future shopping-cart promotions and (2) decreasing the price customers subsequently paid for a product, possibly due to their strategic search for lower prices. Importantly, these long-term effects of price promotions on consumer engagement and strategic behavior spilled over to sellers that did not previously offer promotions to customers. Heterogeneous treatment effects across promotion, seller, and consumer characteristics are examined. This research documents the causal effects of dynamic pricing through price promotions on consumer behavior on a retailing platform, which have important implications for platforms and retailers.

Key words: dynamic pricing, price promotions, platform operations, retail operations, field experiment, empirical OM

---

1. Introduction

With the explosion of available data about customers' purchases, browsing habits, and mobile locations, a growing number of retailers have begun to use dynamic pricing as a revenue management

* The first two authors contribute equally.

tool (Özer et al. 2012). The essence of dynamic pricing involves changing prices based on algorithms that take into account internal factors (e.g., inventory, shipping terms) and external factors (e.g., competitor pricing, market demand). One prevalent way that retailers implement dynamic pricing algorithms is through offering price promotions (Natter et al. 2007, Wu et al. 2014, Ferreira et al. 2015, Cohen et al. 2017). For example, Amazon, which is famous for its implementation of dynamic pricing, frequently offers all kinds of promotions (e.g., Deal of the Day, Lightning Deals, limited-time sales). Abundant decision support systems (e.g., Competera, Market Track) enable retailers to adjust prices and offer promotions more frequently than would humans ever be able to do and have been widely adopted by retailers of all sizes[1].

Besides retailers' growing interest in dynamically adjusting prices through promotions, another trend in the retailing industry is the fast development of two-sided online marketplaces where many retailers sell products on the same platform (e.g., Amazon, Alibaba, JD.com, Groupon). These online retailing platforms further enable retailers to capitalize on various data sources (e.g., demand, consumer characteristics, consumer ratings, and competitor pricing) to develop dynamic pricing and promotion strategies. Retailers extensively leverage the breadth of the data made available by platforms, and dynamically change price through offering promotions to stay competitive.

Considering these trends in the retailing industry, it is critical to understand (a) how dynamic pricing on retailing platforms, especially through offering promotions, influences customer behaviors both in the short and long term, as well as (b) how these effects spill over from retailers that implement a certain dynamic pricing algorithm to other retailers on the same platform. Despite their importance, these questions have not been widely studied empirically due to a lack of appropriate field data from platforms and the difficulty of providing causal inferences. We examine these issues by conducting a large-scale field experiment with Alibaba Group, the world's largest retailing platform.

In particular, we study a promotional tool designed by Alibaba that enables retailers to give coupons to customers who have already added promoted products to their shopping carts but have not purchased them within 24 hours. From the retailer's perspective, this promotional tool allows them to not only frequently adjust their price to manage their sales, but also implement price discrimination by focusing on consumers who have promoted products in their carts and further selecting the customer segment they want to target (e.g., based on demographics, geo-location). From the customer's perspective, if a retailer uses this promotional tool, the price of a product could change along the lifetime of the product in response to consumers' specific actions.

[1] http://www.capterra.com/pricing-optimization-software/;
https://blog.statsbot.co/strategic-pricing-d1adcc2e0fd6;
https://retail.emarketer.com/article/retailers-change-pricing-models-avoid-race-bottom/58dec16debd4000a54864adf

Measuring the effects of a price promotion is complicated: simply comparing outcomes between customers experiencing a price promotion and customers not experiencing one is inappropriate because of the selection inherent in the process of determining who should receive a promotion at which point in time. In our case, an estimate of such a simple comparison may be driven by the inherent differences between customers who are willing to add items to their carts without purchasing them in 24 hours and customers who do not tend to keep items in their carts. To causally identify the short- and long-term effects of our promotion program, we conducted a randomized field experiment with more than 100 million customers and $11,000$ retailers. Among all customers who were eligible for our program, we randomly assigned half of them to receive coupons (treated customers): these customers received coupons if they added promoted products to their carts and met additional criteria (which we describe in Section 4.1). The other half were withdrawn from receiving coupons (control customers): these customers did not receive coupons even if they added promoted products to their carts and met all criteria. Our treatment period was from March 12, 2016 to April 11, 2016, and we continued to observe customers' search and purchasing activity during the first, sixth and twelfth months following our treatment period.

We focus on five main questions in this paper. First, we quantify the effectiveness of our promotion program in the short term. We find that being exposed to this promotion program (i.e., the intent-to-treat effect) boosted customers' purchasing probability and the revenue (net the cost of promotion) of promoted products by 116% and 90%, respectively.[2] Moreover, using an instrumental variable analysis, we demonstrate that viewing the coupons increased customers' purchasing probability of the promoted products by 173%.

Second, and more importantly, we test whether our promotion program goes beyond its role as a short-term price-reduction tool and attracts customers to become more engaged with the platform in the long term. We conduct an intent-to-treat analysis to estimate the long-term effects of being exposed to our promotion program on consumer engagement. We find that being exposed to shopping-cart promotions increased the number of products customers viewed per day on the platform by 0.53% and the likelihood of purchasing any product on a given day by 0.29% during the month following our treatment period. We also use an instrumental variable approach to causally estimate that viewing shopping-cart coupons in the treatment period subsequently increased the number of daily product views by 0.68% and purchase incidence by 0.37%. Furthermore, we find that these effects persisted six months after our treatment period, but wore off in twelve months.

[2] We use "exposure" to refer to the incident where a price promotion is offered to a customer. Being exposed to a promotion does not mean that the customer views the promotion or uses the promotion to purchase a promoted product. Whether or not a customer is exposed to our promotion program depends on whether the customer is randomly assigned to the treatment or control group. Thus, we refer to the effect of being exposed to our promotion program as the intent-to-treat effect.

Third, we assess how dynamic pricing changes customers' strategic behavior in the long term. We show that our promotion program made customers more strategic in two ways. First, our promotion program caused customers to add a greater proportion of products they had viewed to the shopping cart in the month following the treatment period. Specifically, being exposed to shopping-cart promotions and viewing coupons in the treatment period subsequently increased the ratio of products being added to the cart by 0.41% and 0.52%, respectively. These findings suggest that our treatment led customers to strategically add products to the cart to trigger the mechanism via which they previously obtained coupons. Second, we show that customers who were affected by our promotion program paid a lower price than control customers for identical products in the one-month post-treatment period. Specifically, being exposed to shopping-cart promotions and viewing coupons in the treatment period subsequently decreased the price paid for a product that did not offer shopping-cart promotions by 0.02% and 0.03%, respectively. This suggests that customers who were affected by our promotion program managed to purchase the same products at a lower price in ways other than getting shopping-cart coupons (e.g., participating in other promotional activities or waiting for the listing price to go down). In other words, receiving initial promotions led customers to have a lower willingness to pay and engage more in strategic searching for price-reduction mechanisms other than the shopping-cart promotions.

Fourth, we examine whether consumer interaction with retailers who did not previously offer shopping-cart promotions to customers (hereafter, "no-promotion sellers") is also affected by consumers' prior exposure to shopping-cart promotions. Interestingly, we find that the positive long-term effects of receiving and viewing shopping-cart promotions on customers' search and purchasing activity spilled over to retailers that did not previously offer shopping-cart promotions. In a similar vein, we find that the increase in strategic searching and cart add-ons also spilled over to no-promotion sellers.

Last, we demonstrate how the impact of promotions on customers' short- and long-term behaviors is moderated by various characteristics of promotions, customers, and promotion sellers. We find that our promotion program in general changed consumer behavior more effectively in the short or long term when (1) promotions were offered for products that did not belong to promotion sellers' main industry, (2) promotion depth was higher, (3) promotion sellers had a larger gross merchandise value (GMV), (4) consumers were more experienced and (5) customer received promotion more frequently. In addition, we explore which type of seller is more likely to be affected by consumers' exposure to promotions in the long term. We have suggestive evidence that sellers with a large GMV are affected more than those with a small GMV. These results about the heterogeneous treatment effects of promotions can be used to better design promotions and target consumers so as to balance the positive effects of promotions on short-term sales and long-term engagement with the negative effects on long-term strategic behavior.

2. Literature Review and Major Contributions

Our research is connected to and contributes to four streams of literature. First, our research is related to a large literature on dynamic pricing, often focusing on building pricing models and algorithms in various industries to maximize profits (see Özer et al. (2012) for a detailed review). Recently, an emergent literature about dynamic pricing has focused on measuring the effectiveness and consequences of dynamic pricing empirically. Li et al. (2016) documented large operational and financial benefits of adopting dynamic pricing in a "sharing economy" platform. Ferreira et al. (2015), Cheung et al. (2017), and Xu et al. (2016) proposed new dynamic pricing algorithms, such as through price promotions, implemented such algorithms in practice, and measured their effectiveness. Cohen et al. (2017) solves the dynamic promotion planning problem using linear programming and uses supermarket data to measure the effectiveness of their algorithms.

We contribute to this literature in three ways. First, we investigate how implementing dynamic pricing through price promotions affects customer behavior on an online retailing platform, in contrast to past empirical Operations Management literature that has focused on short-term revenue lifts brought by a price promotion policy. Second, unlike past literature on dynamic pricing examining one retailer that implements a specific dynamic pricing algorithm, our field experiment involves more than $11,000$ sellers on a platform, allowing us to measure the spillover effects of a price promotion policy to sellers that did not previously engage in such a pricing activity on the platform. Third, the extant literature has *assumed* that frequent price changes, such as offering promotions, are very costly to retailers and thus have investigated algorithms with limited price changes (Netessine 2006, Broder 2011, Cheung et al. 2017). In this paper, we empirically demonstrate both the positive and negative effects of price changes on customers' long-term behavior, which justifies as well as improves on this assumption in the literature.

The second related stream of research is the operations management literature on strategic customer behavior. This literature has largely focused on two issues. The first focus is to document different ways in which customers exhibit strategic behaviors. For example, Li et al. (2014) demonstrate that customers may expect a future discount and thus strategically time their purchases to obtain lower prices in the airline industry. Moon et al. (2017) show that customers are more likely to be strategic once the retailer reduces the costs of obtaining information about products' prices and availability. The second focus of this literature is on using analytical models to solve traditional operations problems, such as pricing (Cachon and Feldman 2015), inventory planning (Su 2007, Su and Zhang 2008), and capacity management (Allon and Zhang 2015), in the presence of strategic customers. All of these studies have treated strategic customer behavior as an exogenous factor and assumed that at least a fraction of customers are strategic. However, we treat strategic customer behavior as endogenous and examine how people become more strategic after experiencing price promotions.

We also contribute to an emergent literature that studies operations problems on platforms (Kabra et al. 2015, Li et al. 2016, Cui et al. 2016, Taylor 2016, Tang et al. 2016). A large body of this literature has focused on how to match supply with demand using dynamic pricing as a tool (Cachon and Feldman 2015, Tang et al. 2016), assuming that agents on both sides are rational and forward looking. Cachon and Feldman (2015) study the efficiency of dynamic pricing and fixed pricing on platforms and demonstrate that dynamic pricing with fixed ratios leads to near-optimal profits and consumer welfare. Taylor (2016) combines a two-sided platform model with a queuing model and characterizes the impact of delay sensitivity on optimal prices. Most papers in this literature on platform pricing are analytical, with an exception of Li et al. (2016), which provides empirical evidence that professional service providers significantly outperform non-professionals due to their ability to dynamically change prices. We contribute to this literature by empirically examining the spillover effects of dynamic pricing through price promotions across sellers on a platform.

Lastly, our research substantially extends the marketing literature about the long-term effects of price promotions. This research stream has primarily relied on UPC scanner data to estimate how in-store price promotions change consumers' purchasing activities over an extended period (e.g., Mela et al. 1997, 1998, Ailawadi and Neslin 1998, Jedidi et al. 1999, Nijs et al. 2001, Pauwels et al. 2002, Anderson and Simester 2004). Our major contribution to this literature lies in our novel research questions and the research context we study. First, extending prior work that examines the long-term effects of promotions on the promoted brand or the corresponding product category, we provide the first investigation of the long-term spillover effect of promotions—that is, how certain retailers' use of price promotions affect other retailers on the same platform. Second, going beyond brand choice and category purchase decisions—the focus of past research—we take a step forward to additionally examine how price promotions change consumers' search behavior in the long term. Third, we find that customers become more strategic after receiving a price promotion by taking action to trigger the same mechanism via which they previously obtained coupons. We further show how such increase in strategic behaviors is moderated by promotion and buyer characteristics, which sheds light on better promotion designs. As far as we know, the previous literature has not studied this mechanism in the field and has not documented how different promotion and consumer characteristics can moderate this mechanism. Lastly, as online retailing platforms become increasingly important in the economy, our study of price promotions on the world's largest retailing platforms is economically meaningful, particularly given that the previous literature has mostly focused on price promotions in brick-and-mortar and mail catalog channels offered by one retailer. In Online Appendix Section 1, we provide more detail about the differences in research questions, methodology (a large-scale randomized field experiment), and context (a large online retailing platform) between our research and the existing marketing literature on the long-term effects of price promotions.

3. Hypothesis Development

We develop three groups of hypotheses. The first group relates to the short-term and direct effects of price promotions on the purchase incidence and revenue of promoted products. The second group relates to the long-term effects of price promotions on customers' engagement on a platform after the promotion period. The last group of hypotheses considers how price promotions alter strategic customer behavior in the long run. For the second and third groups of hypotheses, we also examine the spillover effects of price promotions to retailers that have not offered promotions to customers.

3.1. Short-term effects of promotions

We hypothesize that price promotions increase the purchasing probability of promoted products, assuming that customers' utility from purchases decreases with product prices. When a price promotion is available, a customer can derive a higher utility from purchasing a product and thus should be more likely to purchase it, as compared to the counter-factual case where a price promotion is not available. This hypothesis is straightforward and consistent with past field experiments causally showing that price promotions temporarily increase the purchase likelihood of promoted products in an online environment (e.g., Luo et al. 2014, Fong et al. 2016). However, it is less obvious how price promotions change revenue, where the revenue of each product j is defined as the multiplication of the number of purchased items and the revenue per item: $\text{revenue}_j = \text{number of items purchased}_j \times \text{revenue per item}_j$.

In this case, the effect of price promotions on the revenue of promoted products is unclear, since price promotions may increase the number of purchased items but decrease the revenue per item.[3] In a recent study, Zhang et al. (2017) demonstrate a positive effect of price promotions on the revenue of promoted product. Ferreira et al. (2015) also implement a price reduction algorithm in Rue La La and demonstrate that price reductions lead to higher revenue. Given that we similarly examine price promotions offered in an online retailing environment, we hypothesize that price promotions can increase the revenue of promoted products during the promotion period in our context.

[3] In this paper, "product" refers to a particular SKU offered on the platform, whereas "item" refers to a concrete object. For example, a customer may buy multiple items of the same product at a given transaction.

3.2. Long-term effect of promotions on customer engagement

In this section, we discuss the impact of promotions on customers' long-term engagement, including their views of products and purchases. We employ a theoretic search model to form our hypotheses (Stigler 1961, Burdett and Judd 1983). In such non-sequential search models, an incoming customer decides first whether to visit a platform and then her search behavior on the platform, given her belief about the price distribution of available options. If her expected utility from conducting a search on the platform is lower than her opportunity cost, the customer will simply not visit the platform. Conditional on visiting the platform, the customer chooses the number of products to search. Her objective is to minimize the sum of her total search costs and the expected lowest searched price.

Product Views: In the model above, price reductions through promotions can impact customers' future search behaviors in two ways. First, prior promotions can decrease customers' expectation of prices (Lattin and Bucklin 1989, Kalwani and Yim 1992, Anderson and Simester 2004). Thus, if the optimal number of products to search is held constant conditional on visiting a platform, promotions will increase customers' expected utility from visiting the platform. This will in turn increase customers' probability of visiting the platform and boost the number of products they view on the platform. Second, if consumers notice that only some retailers are using certain promotions, such promotions may increase customers' perceived price dispersion among retailers. Holding the probability of visiting a platform constant, an increase in perceived price dispersion can increase the optimal number of searches that a customer wants to conduct on the platform and in turn boost their product views.

Both of these mechanisms may affect all retailers on the platform, regardless of whether a retailer previously offered a promotion to consumers. In terms of the first mechanism, if consumers expect greater utility from visiting a platform after receiving a promotion there, their increased likelihood of visiting the platform should not only boost their overall product views on the platform but also cause consumers to be more likely to view products offered by no-promotion sellers. In terms of the second mechanism, if consumers perceive a greater price dispersion across retailers on a platform after they receive a promotion from a particular retailer, the increase in the optimal number of searches they want to conduct may cause consumers to view more products in general on the platform, including products from retailers who did not previously offer them promotions. To summarize, we hypothesize that price promotions will increase customers' search activity (e.g., the number of products they browse on a given day) on the platform, which could spill over to consumers' views of products from no-promotion retailers.

Purchase Incidence: It is unclear how price promotions change customers' purchases on the platform in the long run. On one hand, according to the theoretical search model described above, if customers search for more products after receiving price promotions and if the rate at which consumers purchase products upon viewing them remains unchanged, they should be more likely to purchase products on the platform after receiving price promotions. However, a large literature on customer behavior suggests that prices of past purchases can serve as reference points for future purchases (Kahneman and Tversky 1979, Thaler 1985, Lattin and Bucklin 1989, Anderson and Simester 2004, Kőszegi and Rabin 2006, Popescu and Wu 2007). A customer's utility of buying a product depends not only on her intrinsic valuation of the product and the product's actual price (i.e., acquisition utility) but also on the disparity between the actual price and her expectation about the product's price (i.e., transaction utility; Thaler 1985, Lattin and Bucklin 1989, Popescu and Wu 2007, Nasiry and Popescu 2011). At a given actual price level, as reference price reduces, transaction utility decreases and so too does purchase likelihood. In this vein, receiving a price promotion could decrease a customer's reference price, increase the disparity between regular prices and her reference point in the long run (Lattin and Bucklin 1989, Kalwani and Yim 1992), and consequently decrease her purchasing probability. Combining all arguments, we reach competing hypotheses about whether price promotions increase, decrease, or have no effect on customers' purchase likelihood of products from promotion sellers as well as non-promotion sellers on the platform in the long run. The long-term impact of our price promotion program on purchase likelihood depends on an interplay of competing forces: increased search activity versus decreased reference price.

Expenditure: Similarly, it is unclear how price promotions change customers' expenditures on the platform in the long run. On one hand, after price promotions, increased search activity may increase purchase incidence and lead customers to spend more money on the platform. However, as explained above, increased search activity may decrease purchase incidence due to decreased reference prices. Even if purchase likelihood is boosted, overall expenditures may not increase because price promotions may lead customers to become more strategic about timing their purchase and buy products at lower prices, which we elaborate on in Section 3.3. Thus, we have competing predictions about how price promotions change customers' expenditures on the platform.

3.3. Long-term effects of promotions on customers' strategic actions

In this section, we propose that receiving price promotions may train consumers to become more strategic in two ways and that such effects can spill over to no-promotion retailers. First, receiving a price promotion via a certain mechanism may lead consumers to expect that more of such promotions will be available from the same mechanism in the future. As a result, people may be tempted to take action to trigger the mechanism via which they previously received a promotion. In our context, after receiving a price promotion targeted at products in their shopping carts, consumers may expect future promotions to be available for products that they add to the cart. Under this expectation, customers may increase the proportion of items they add to cart after viewing them to increase their odds of receiving shopping-cart-specific price promotions. We denote this as a **direct strategic effect**. If customers are not fully aware of which retailers offer a specific kind of promotion, they may expect many retailers to offer that type of promotion and thus engage in the direct strategic effect even when they interact with no-promotion sellers. In our case, consumers cannot tell whether a retailer offers shopping-cart promotions because the platform does not label retailers differently based on their participation in the shopping-cart promotion program. Thus, we hypothesize that the direct strategic effect will spill over to sellers who did not previously offer shopping-cart promotions, causing consumers to add to cart a larger proportion of products from these sellers.

Second, receiving a price promotion may decrease customers' reference prices and in turn make them search more for lower prices (Thaler 1985, Anderson and Simester 2004, Popescu and Wu 2007). Relatedly, prior work suggests that cumulative exposure to price promotions could increase price sensitivity (Mela et al. 1997, 1998, Jedidi et al. 1999). As a result, customers who previously received a promotion may pay a lower price to buy the same product than customers who did not previously receive a promotion. Importantly, we expect this process to even occur with products that do not offer the same type of promotion as the one that consumers previously received; in other words, we expect consumers who previously received one type of promotion to obtain a lower price via other price-reduction mechanisms. We denote this as a **cross-mechanism strategic effect**. In our case, there are generally two types of price-reduction mechanisms that customers can explore on Alibaba besides shopping-cart promotions. The first price-reduction mechanism involves retailer-specific price promotions. Besides shopping-cart promotions, Alibaba designs other promotional tools for retailers and changes the availability of the tools across retailers over time. For example, during our study period, some retailers have the option to send promotions to their existing consumers via in-app messages, and some retailers can set up a game called "password coupon" such that consumers can search for a password on Alibaba and receive a discount upon finding the password. Retailers could also announce their promotion dates for certain subsets of products, and such dates are normally determined by retailers' product launch schedule or special anniversaries. Strategic consumers can participate in these promotional activities or time their purchase to wait for a retailer's promotion date. The second mechanism involves platform-wide price promotions. Alibaba has several special "holidays" when it gives out platform-wide price reductions. The most notable one is "double eleven" (November 11), a 24-hour shopping day when all retailers have to offer prices at least 10% lower than their lowest prices of that year. Consumers who want to obtain a low price can wait for platform-wise promotions. To summarize, we hypothesize that

Figure 1 An example of the shopping-cart promotion

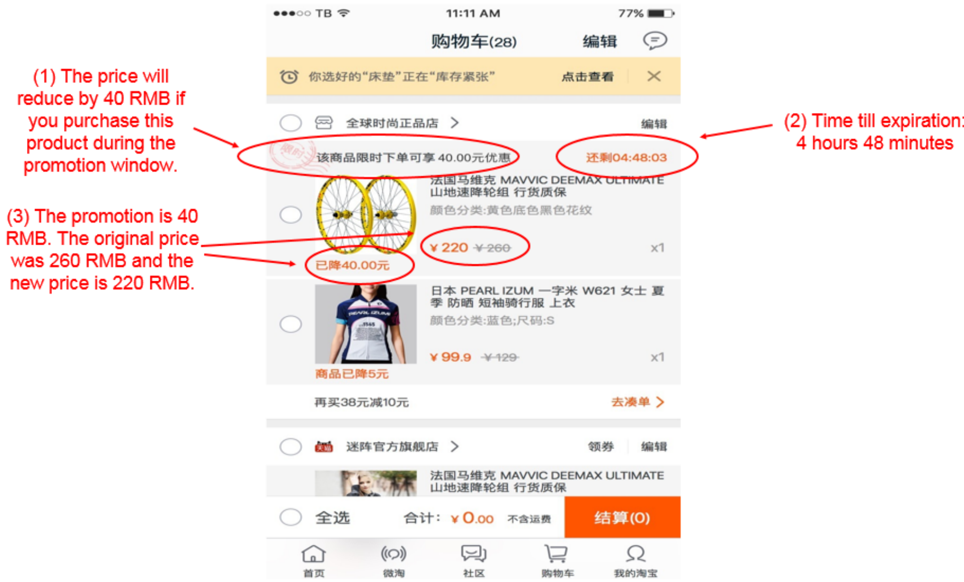

receiving price promotions leads customers to pay a lower price for the same product than would they if they were not exposed to price promotions, even for products that do not offer the same type of promotion as what they previously received. Furthermore, our hypothesized pattern may spill over to products from retailers who did not previously offer shopping-cart promotions, because prior promotions may reduce consumers' reference prices and increase their price sensitivity when they shop on the platform in general even when they purchase from no-promotion sellers.

4. Randomized Field Experiment Overview and Data

4.1. Field Setting and Experimental Design

We conducted a large-scale longitudinal field experiment in collaboration with Alibaba Group, a Chinese online and mobile commerce company. Alibaba was founded in 1999. As of 2016, its Gross Merchandise Value (GMV) had surpassed 3 trillion RMB (equivalent to 485 billion US dollars).[4] This makes Alibaba the world's largest retailer, overtaking Walmart, which posted revenues of $482.1 billion for the same year.[5] The randomized field experiment was run on Taobao marketplace, China's largest peer-to-peer retailing platforms for small businesses and individual entrepreneurs, as well as on Tmall.com, China's largest third-party business-to-consumer platform for branded goods. For the sake of simplicity, we hereafter refer to Taobao Marketplace and Tmall.com jointly as "the Alibaba platform."

[4] http://www.alibabagroup.com/en/news/press_pdf/p160505.pdf

[5] https://www.forbes.com/sites/jlim/2016/05/05/alibaba-fy2016-revenue-jumps-33-ebitda-up-28/5f15a05f53b2

Figure 2 Study Flow from the Customer's Perspective

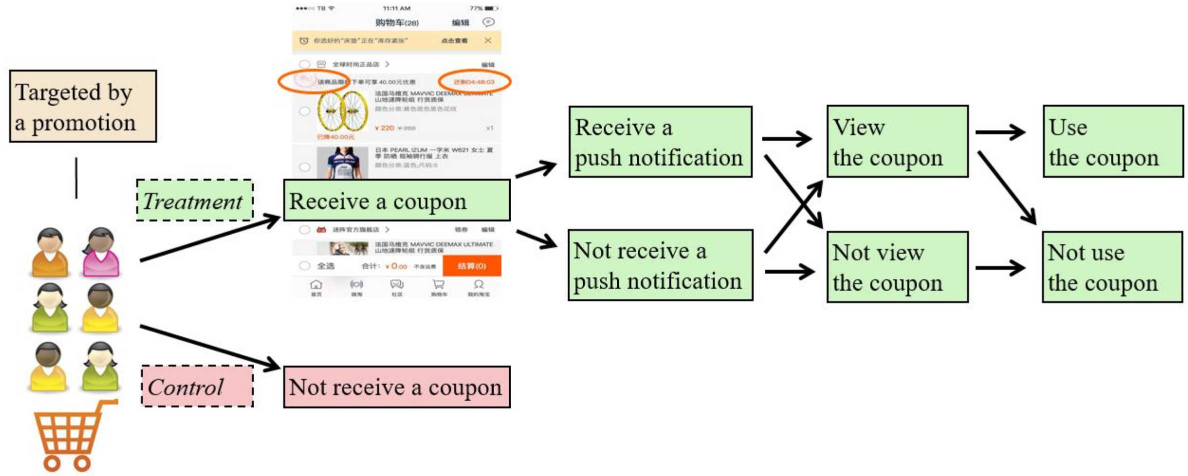

On March 12, 2016, the Alibaba platform officially introduced its "Shopping-cart Promotion Program" (hereafter, "the promotion program"), which allows sellers to send time-limited coupons to selected customers who have already added the sellers' products to their shopping carts without purchasing them in 24 hours.[6] To launch a price promotion, a seller first selects the product that she wants to promote, her promotion price, and the customer segment that she plans to target (based on characteristics such as customers' demographics, geo-location, and tenure). The size of the discount has to be larger than 5 RMB. Within the pool of targeted customers, customers are randomly assigned to a treatment group that actually receives the promotion or a control group that is excluded from receiving the promotion. Random assignment is determined based on the last digit of customers' Alibaba ID numbers, which were generated randomly. During our study period, customers assigned to one group remained in the same group.

To summarize, for each customer-product pair, the customer is targeted by a shopping-cart promotion if the following three criteria are satisfied: (1) The seller of the product has chosen to run the promotion program on the product; (2) The customer has added the product to her cart for more than one day without making a purchase; (3) The customer's characteristics are selected by the seller. For a targeted customer to receive the promotion, she has to satisfy an additional criterion: (4) The customer has to be in the treatment group based on her Alibaba ID.

[6] The Alibaba platform focuses on promoting products that are already added to the cart for two reasons. First, adding products to the cart is the last step before consumers make purchases. Sending consumers coupons at this stage of online shopping may be more effective than sending consumers coupons at an earlier stage (e.g., while consumers are browsing products). Second, shopping cart abandonment, whereby consumers add products to the shopping cart without making a purchase, is a challenge faced by the online retail industry. According to statistics from Alibaba, the rate at which consumers in this program purchase products in their shopping carts drops sharply from an average of 16.79% on the day when they add the products to the cart to 2.82% on the next day.

Figure 3 Timeline of Our Data Sample

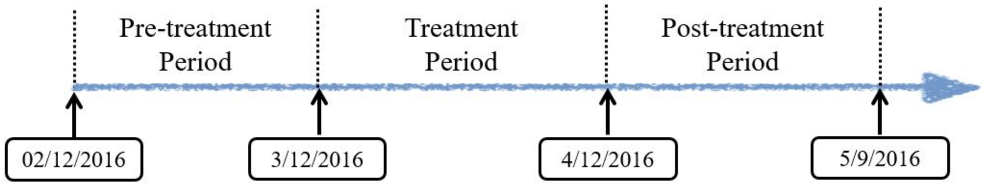

The promotion program is only available on the mobile app of the platform.[7] If a customer is targeted by a promotion and is in the treatment group, a message shows up above the promoted product in the shopping cart in her mobile app. Figure 1 shows an example of such a message. After clicking on her shopping cart in the app, a customer will see (1) the promotion, (2) the size of the discount, and (3) the remaining time before the coupon expires. To use a coupon from the promotion program, customers need to purchase the promoted product in the app before the coupon expires at the end of the day. As discussed, if a customer is targeted by a promotion but is in the control group, she does not receive the coupon. Figure 2 depicts the design of the experiment from the customer's perspective, in terms of the experiences and actions of customers in the treatment and control groups.

4.2. Data and Randomization Check

We randomly drew *1 million* customers who were targeted by at least one promotion between March 12, 2016 and April 11, 2016: 500,000 in the treatment condition and 500,000 in the control condition.[8] We tracked these customers' search and purchasing activities across all sellers on the platform from February 12, 2016 to May 9, 2016. During this observation period, no other large promotions targeted only the treatment or control group of customers.

The data sample in our main analysis can be divided into three time periods. The first period is the *pre-treatment period* from February 12, 2016 to March 11, 2016. The data in this period is used to confirm that treated and control customers are comparable prior to our experiment, which we should expect by random assignment. The second period is the *treatment period* from March 12, 2016 to April 11, 2016. In this period, customers in the treatment group received coupons, while customers in the control group did not. We use the data in this period to estimate the short-term effectiveness of our promotion program. The last period is the *post-treatment period* from April 12, 2016 to May 9, 2016. We use the data from this period to measure the long-term impact of promotions on customer behavior. Figure 3 demonstrates the timeline of the data sample for our

[7] Majority of Alibaba consumers use its mobile app. According to Alibaba's public report (http://www.alibabagroup.com/en/news/press_pdf/p160505.pdf), in Q1 of 2016 (i.e., our experiment period), mobile GMV accounted for 73% of total GMV. Annual active buyers on the Alibaba platform reached 423 million in March 2016, while mobile annual active users were 410 million at that time.

[8] We are limited to 1 million customers due to security constraints from Alibaba Group.

analysis reported in Sections 5, 6, and 7. In addition, we observe consumer behavior during the sixth and twelfth months following our treatment period, which we use to explore how long-lasting the effects of our promotion program were (see Section 8.1 for detail).

The original price of promoted products ranged from 6 to $11,110$ RMB in our sample (equivalent to approximately 1 to $1,760$ US dollars; Mean = 120 RMB). Promotion depth (discount divided by original price) ranged from 8% to 99%, and the average promotion depth of 17% suggests that sellers in our program generally offer moderate discounts (as opposed to deep discounts like those available on Groupon). The gender (47% females among those with known gender information) and age (mean = 29, median = 27) of consumers in our sample are comparable to the characteristics of the population of consumers who use Alibaba's mobile app (51% females, 57% at age 29 or younger)[9]. On average, customers in the treatment condition received shopping-cart promotions 1.3 times during our treatment period, with 80% of treated customers receiving only one promotion.[10] In our treatment period, promotion sellers came from 21 product categories, with 55% being Tmall sellers and 45% being Taobao sellers.

Following Alibaba's guidelines, in this paper, we exclude customers who purchased more than 500 products or viewed more than $50,000$ products on the platform on any given day during our observation period, because these customers were likely to shop for organizations rather than households. These customers account for 6.9% of the $1,000,000$ customers we randomly drew. Our results are robust if we include these customers. The size of the raw data is 500 gigabytes which is aggregated from more than 1 petabytes of data in Alibaba's internal database. Out of computation feasibility, for some analyses, we have to aggregate data to the customer-day level to compose outcome measures (e.g., daily views of products, daily expenditure).

Due to the large sample size of our experiment, we are concerned that we may capture random statistical fluctuations between the treatment and control groups that are not caused by our experimental treatment (i.e., the offering of shopping-cart promotions). Therefore, we first use the pre-treatment data (i.e., data from February 12, 2016 to March 11, 2016) to check whether our treatment and control groups are statistically different in any important metrics prior to our experiment. For each day in the pre-treatment period, we calculate the number of products that customers viewed, the number of products they added to their shopping carts, and their expenditures. As shown in Table 1, despite of our large sample size, customers in the treatment group

[9] http://www.alibabagroup.com/en/ir/pdf/160614/02a.pdf

[10] A consumer theoretically could still be targeted by shopping-cart promotions in the subsequent months following our treatment period. However, the design of the shopping-cart promotion program makes it difficult for the same consumer to be targeted by the program more than once given the limited number of sellers offering this promotion. In fact, Alibaba's internal statistics show that nearly 80% of consumers who were ever targeted by the shopping-cart program were targeted only once before the program discontinued. Our results remain qualitatively the same even if we only focus on customers who were only targeted once.

Table 1 Baseline Characteristics and Randomization Checks

	Daily Views	Daily Number of Products Added to Cart	Daily Expenditure (RMB)
Treatment Group	21.21	1.58	23.79
Control Group	21.21	1.58	23.78
Difference	0.00	0.00	0.01
T-test p-value	0.5711	0.6083	0.9433

Note: This table reports the mean daily values of three variables across all customers in the treatment vs. control groups during the pre-treatment period (February 12, 2016 - March 11, 2016). T-tests are based on 26,999,261 observations (i.e., $931,009$ customers $\times$ 29 days $= 26,999,261$ consumer-day pairs).

have statistically indistinguishable metrics as compared to those in the control group: on average, they similarly viewed 21.2 products on the app, added 1.6 products to their carts, and spent approximately 23.8 RMB each day. We further confirm that our random assignment is successful by conducting an ordinary least squares regression on consumers' pre-treatment daily search and purchasing activity and controlling for date fixed effects (Section 2 in Online Appendix). The results of our randomization checks suggest that any difference after the promotion program was implemented should be attributed to the program.

5. The Short-term Effects of Promotions

We first examine the short-term effects of the promotion program on promoted products. Our unit of analysis is each customer-product pair ijt in which customer i was targeted by a price promotion for product j on day t in the treatment period, and we have $1,693,239$ observations. In Section 5.1, we first examine how being exposed to a shopping-cart promotion affects customers' decisions about purchasing a product, denoted as the **intent-to-treat effect of promotions**. In Section 5.2, we then estimate the effect of seeing a promotion on customers' purchasing decisions, denoted as the **average treatment effect of viewing promotions**.

5.1. Short-term Intent-to-treat Effects of Promotions

We start with an intent-to-treat analysis that simply compares purchasing behavior between customers in the treatment and control conditions. Panel A of Table 2 shows that the average probability of customers purchasing a promoted product more than doubled from 0.98% in the control condition to 2.11% in the treatment condition; the average expenditure on a promoted product increased by 91% from 1.01 RMB in the control condition to 1.93 RMB in the treatment condition. T-tests and non-parametric Wilcoxon tests both confirm the statistical significance of these comparisons (all p-values < 0.0001).

Table 2 Short-term Effects of Promotions

Panel A: Summary Statistics		
	P.I. (%)	Expenditure
Treatment Group	2.11%	1.93
Control Group	0.98%	1.01
Difference	1.13%	0.92
T-test p-value	< 0.0001	< 0.0001
Wilcoxon p-value	< 0.0001	< 0.0001

Panel B: OLS Regression Results				
	Intent-to-treat Effects		Average Treatment Effects	
	P.I.	Expenditure	P.I.	Expenditure
	(1)	(2)	(3)	(4)
Treatment	0.0114****	0.9146****		
	(0.0002)	(0.0248)		
Viewing Coupon			0.0170****	1.3678****
			(0.0003)	(0.0370)
Relative Effect Size	116%	90%	173%	135%
Product Fixed Effects	Yes	Yes	Yes	Yes
Date Fixed Effects	Yes	Yes	Yes	Yes
Observations	1,693,239	1,693,239	1,693,239	1,693,239

Note: $*p < 0.10; **p < 0.05; ***p < 0.01; ****p < 0.001$. Robust standard errors are in parentheses.

Note: Panel A reports the average outcomes of promoted products across subjects in the treatment and control groups during the treatment period (March 12, 2016 - April 11, 2016). In Panel B, Columns (1) and (2) report the results from OLS regressions that estimate the intent-to-treat effects of being exposed to a promotion due to the treatment assignment. Columns (3) and (4) report the results from OLS regressions that estimate the average treatment effect of viewing a coupon using instrumental variable (IV). For each outcome variable, the relative effect size is computed as the coefficient divided by the average value in the control group. In all tables in this paper, "P.I." standards for "Purchase Incidence."

We next conduct regression analyses to determine whether the effects described above are robust to controlling for other variables that may affect purchasing behavior.[11] Specifically, we use the following OLS regression specification:

$$\text{Purchase Incidence}_{ijt} = \alpha_0^1 + \alpha_1^1 \text{Treatment}_i + X_j + D_t + \epsilon_{ijt} \tag{1}$$

$$\text{Expenditure}_{ijt} = \alpha_0^2 + \alpha_1^2 \text{Treatment}_i + X_j + D_t + \epsilon_{ijt} \tag{2}$$

where Purchase Incidence_{ijt} is a binary variable that equals 1 if customer i purchased promoted product j on day t, and 0 otherwise; Expenditure_{ijt} is the total amount of money in RMB that customer i spent on product j on day t, which equals 0 if consumer i did not purchase product j; Treatment_i is a binary variable indicating whether customer i is in the treatment group. Table 7 in the Appendix summarizes and defines the key dependent and independent variables examined in this paper. X_j represents product fixed effects, and D_t represents date fixed effects. We report

[11] This is a field experiment with proper randomization and a large sample size. Therefore, the estimators should be unbiased without controls and control variables are added only to make the estimators more efficient.

robust standard errors in this analysis as well as all subsequent analyses presented in this paper. However, our results all hold true if we cluster standard errors at the customer level or employ a logistic regression on Purchase Incidence_{ijt} (a binary dependent variable).

Columns 1 and 2 in Panel B of Table 2 present results from Specifications (1)-(2). The positive and significant coefficients on treatment indicate that, as compared to the control condition, receiving coupons significantly increased the purchase probability of promoted products by 116% and the money spent by 90% (both p-values < 0.001).

5.2. Short-term Average Treatment Effects of Viewing Promotions

So far, we have shown that being in the treatment group can significantly lift the purchase incidence and revenue of promoted products. However, customers in the treatment condition may or may not notice the coupons sent to their mobile app. Customers could not be influenced by coupons if they had not noticed them, and only 66% of customers in the treatment condition actually viewed their coupons in our data. Therefore, to determine the average treatment effect of actually viewing a coupon (instead of the effect of receiving a coupon) on customers' short-term behaviors, we examine the effect of viewing a shopping-cart coupon on customers' purchasing behaviors.

Notice that we cannot estimate this effect causally by simply comparing customers who viewed coupons with customers in the control condition because of omitted variable biases. Following the same procedure used in past literature (Angrist and Pischke 2008, Zhang et al. 2017), we use the random assignment of customers into the treatment and control groups as an instrumental variable (IV) for viewing coupons. We denote whether or not customer i viewed the coupon of promoted product j on day t as Viewing Coupon_{ijt}. For this setup to be valid, we need to check two assumptions of our instruments.

First, we need to satisfy the inclusion restriction assumption: our random assignment should be correlated with whether customers viewed coupons. In other words, this assumption can be summarized as follows:

ASSUMPTION 1. **Inclusion restriction**: $cov(\textit{Viewing Coupon}_{ijt}, \textit{Treatment}_i) \neq 0 \; \forall i, j, t.$

This assumption must be true according to the nature of our experiment: only customers who were randomly assigned to the treatment condition could view coupons. In fact, as described earlier, about 66% of customers in the treatment group viewed the shopping-cart coupon in their mobile app, while 0% of customers in the control group viewed any shopping-cart coupons because they did not receive any.

Second, we need to satisfy the exclusion restriction assumption: that is, random assignment into treatment (vs. control) should only affect dependent variables by making coupons more visible to customers, and should be not correlated with any other uncontrolled variables that affect dependent variables directly (e.g., customers' inherent interest in the promoted product). This assumption can be summarized as follows:

Assumption 2. **Exclusion restriction**:

$$cov(Outcome\ Variables_{ijt}, Treatment_i | Viewing\ Coupon_{ijt}) = 0\ \forall i, j, t,$$

where $Outcome\ Variables_{ijt} \in \{Purchase\ Incidence_{ijt}, Expenditure_{ijt}\}$.

This assumption is satisfied because (a) customers were randomly assigned into the treatment or control group and (b) customers in the two groups only differed in their eligibility for receiving shopping-cart coupons and did not encounter other differential treatment by sellers or the Alibaba platform. Our randomization check in Section 4.2 further confirms that customers in treatment and control groups were indifferent prior to the treatment. Note that coupons were sent automatically to customers and that sellers could not include additional information to market their products. Thus, it is not possible for customers in the treatment condition to receive additional marketing messages. Furthermore, sellers were not aware of this experiment or group assignment and thus could not strategically send different messages to two groups of customers via other channels.

Having confirmed the validity of these two assumptions, we can use the standard IV setup to estimate the average treatment effects of viewing coupons on purchasing behaviors (Angrist and Pischke 2008). Columns 3 and 4 in Panel B of Table 2 shows the estimated results from regressions that control for product and date fixed effects. We find that viewing coupons significantly increased the purchase probability of promoted products by 173% and expenditures by 135% (both p-values < 0.001), relative to the control group. Therefore, we observe that viewing coupons significantly increased the sales of promoted products in the short term.

Altogether, we provide strong evidence that sending targeted shopping-cart promotions effectively elicited customers' immediate responses to promoted products, especially if customers actually viewed the coupons. In the short term, our promotion program is a very effective tool for lifting sales. In the following sections, we answer the important question of how this promotion program influences customers' search and purchasing activity in the long term on the entire platform as well as customers' long-term strategic levels.

6. Long-term Effects of Promotions

6.1. Long-term Effects of Promotions on Customer Engagement

To assess the long-term effects of the promotion program on customers' engagement with the platform, we analyze customers' search and purchasing activity during the post-treatment period (i.e., April 12, 2016 to May 9, 2016). As discussed in Section 3.2, for each customer i on each day t, we consider three variables tracking her engagement on the platform: (1) $Daily\ Views_{it}$: the number of products customer i viewed on day t; (2) *Daily Purchase Incidence*$_{it}$: whether customer i on day t made any purchase on the platform; and (3) *Daily Expenditure*$_{it}$: the amount of money (in RMB) customer i spent on the platform on day t.

We first calculate these variables across all sellers on the platform for each customer i on each day t. We are also interested in whether the promotion program has a spillover effect to sellers who did not previously offer promotions to the customer. Therefore, for customer i on day t, we identify sellers who had not previously targeted customer i and compute the engagement variables associated with these no-promotion sellers. We denote these variables from no-promotion sellers as $\{$*Daily Views*$^o_{it}$, *Daily Purchase Incidence*$^o_{it}$, *Daily Expenditure*$^o_{it}\}$.

We start with an intent-to-treat analysis to directly compare customers in the treatment condition with those in the control condition. This analysis estimates the causal effect of being exposed to a shopping-cart promotion on customer engagement in the post-treatment period. Specifically, we compare the search and purchasing activity of treatment and control customers using the following system of regression specifications:

$$\text{Outcome Variable}_{it} = \beta_0 + \beta_1 \text{Treatment}_i + D_t + \epsilon_{it} \tag{3}$$

$$\text{Outcome Variable}^o_{it} = \beta_0 + \beta_1 \text{Treatment}_i + D_t + \epsilon_{it} \tag{4}$$

where $\text{Outcome Variable}_{it} \in \{\text{Daily Views}_{it}, \text{ Daily Purchase Incidence}_{it}, \text{ Daily Expenditure}_{it}\}$, $\text{Outcome Variable}^o_{it} \in \{\text{Daily Views}^o_{it}, \text{ Daily Purchase Incidence}^o_{it}, \text{ Daily Expenditure}^o_{it}\}$, Treatment_i is a binary variable indicating whether or not customer i is in the treatment group, and D_t represents date fixed effects. We have a balanced panel including $25,137,243$ observations ($931,009$ customers $\times$ 27 days), where each observation is a customer-day pair.

Table 3's Panel A shows the results from Specification (3)-(4). Columns 1-3 include all sellers on the platform. The positive and significant coefficients on treatment in Columns 1 and 2 indicate that customers in the treatment condition viewed more products and were more likely to make any purchase on a day during the post-treatment period compared to those in the control condition. Specifically, the daily number of products viewed increased by 0.53%, and purchase incidence increased by 0.32% among treated customers, relative to the control condition (i.e., 24.47 views per day and a daily purchase rate of 18.81 percentage points on average). In Column 3, the coefficient on treatment is negative and insignificant, suggesting that daily expenditure did not significantly differ between the treatment and control customers during the post-treatment period. Columns 4 to 6 show that the results for no-promotion sellers are almost identical to those for all sellers because no-promotion sellers account for 98% of all sellers in the data sample.

Next, we examine the long-term effects of viewing a coupon on customers' engagement. As explained in Section 5.2, we need to use the random assignment of customers into the

Table 3 Long-term Effects of Promotions on Customer Engagement

Panel A: Long-term Intent-to-Treat Effects						
	Daily (Views	P.I. (All Sellers)	Expenditure)	Daily (Views	P.I. (No-promotion Sellers)	Expenditure)
	(1)	(2)	(3)	(4)	(5)	(6)
Treatment	0.1302****	0.0006****	-0.0116	0.1287****	0.0005****	-0.0116
	(0.0180)	(0.0002)	(0.0769)	(0.0180)	(0.0002)	(0.0769)
Relative Effect Size	0.53%	0.32%		0.53%	0.27%	
Panel B: Long-term Average Treatment Effects of Viewing Coupons						
	Daily (Views	P.I. (All Sellers)	Expenditure)	Daily (Views	P.I. (No-promotion Sellers)	Expenditure)
	(1)	(2)	(3)	(4)	(5)	(6)
Ever Viewing Coupon	0.1670****	0.0007****	-0.0149	0.1649****	0.0007****	-0.0149
	(0.0231)	(0.0002)	(0.0986)	(0.0230)	(0.0002)	(0.0985)
Relative Effect Size	0.68%	0.37%		0.68%	0.37%	
Date Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Observations	25,137,243	25,137,243	25,137,243	25,137,243	25,137,243	25,137,243

Note: $*p < 0.10; **p < 0.05; ***p < 0.01; ****p < 0.001$. Robust standard errors are in parentheses.

Note: This table shows the long-term effects of being in the treatment group (Panel A) and viewing shopping-cart coupons (Panel B) on the daily number of products viewed (Columns (1)-(2)), the daily probability of making any purchase (Columns (3)-(4)), and daily expenditure (Columns (5)-(6)). Columns (1)-(3) include *all* sellers, and Columns (4)-(6) focus on *sellers who did not previously offer shopping-cart promotions.*

treatment versus control group as an instrumental variable (IV) for viewing coupons. We use Ever Viewing Coupon$_i$ to denote whether or not customer i viewed at least one coupon during our treatment period. Table 4's Panel B reports the causal effects of viewing coupons on customer engagement in the post-treatment period. Relative to the control condition, viewing coupons increased daily views by 0.68% and daily purchase incidence by 0.37% for all sellers (Columns 1 and 2; both p-values < 0.0001). The effects of viewing coupons on daily views and purchase incidence with respect to no-promotion sellers have similar magnitudes (Columns 4 and 5; both p-values < 0.0001). Columns 3 and 6 show that viewing coupons did not statistically significantly change daily expenditure for all sellers nor no-promotion sellers during the post-treatment period.

In conclusion, we find that receiving coupons in the treatment period led customers to view more products and buy products more frequently in the following post-treatment month. This suggests that prior interaction with price promotions increased customers' engagement with the platform in the long run. Importantly, customers' engagement increased even for sellers that did not previously offer them shopping cart promotions. Thus, we show that price promotions have a cross-seller spillover effect on the platform.

6.2. Long-term Effects of Promotions on Strategic Customer Behavior

We next examine the effects of price promotions on the development of strategic customer behavior. As explained in Section 3.3, we seek to identify two forms of strategic behavior in our context: direct strategic effect and cross-mechanism strategic effect. In the following sections, we test whether these two strategic behaviors exist. Similar to our analyses regarding customer engagement, we start with an intent-to-treat analysis to estimate the causal effect of being exposed to a shopping-cart promotion on strategic behavior in the post-treatment period. Next, we use an instrumental variable approach to estimate the long-term effects of customers viewing a coupon on their strategic behavior.

6.2.1. Direct Strategic Effect To estimate the direct strategic effect, for each customer i on each day t between April 12, 2016 and May 9, 2016, we calculate Cart-to-view Ratio_{it} which equals the number of products customer i added to the cart divided by the number of products she viewed on the platform on day t. For each customer-day pair it, we also calculate the ratio exclusively for no-promotion sellers (denoted as Cart-to-view Ratio_{it}^{o}). In our intent-to-treat analysis, we use the following specification to estimate the impact of our promotion treatment on cart-to-view ratio:

$$\text{Cart-to-view Ratio}_{it} = \theta_0^0 + \theta_1^0 \text{Treatment}_i + D_t + \epsilon_{it} \tag{5}$$

$$\text{Cart-to-view Ratio}_{it}^{o} = \theta_0^1 + \theta_1^1 \text{Treatment}_i + D_t + \epsilon_{it} \tag{6}$$

Each observation is a combination of customer i and day t. Note that, due to the coding of the dependent variables, our analysis only includes customer-day pairs in which customer i viewed at least one product on day t; otherwise, the dependent variables would not have a valid value. This leaves us with $15,059,021$ and $15,048,820$ observations for all and no-promotion sellers.

In Columns 1 and 4 of Table 4's Panel A, the positive and significant coefficients on treatment indicate that treated consumers added a greater proportion of viewed products to their shopping cart than control consumers in the post-treatment period (both p-values < 0.0001). Specifically, the cart-to-view ratio in the treatment condition increased by 0.41% for all sellers and 0.41% for no-promotion sellers, relative to the control condition (i.e., an average ratio of 7.637 and 7.644 percentage points for all sellers and no-promotion sellers, respectively). Columns 1 and 4 of Table 4's Panel B reports the causal effects of viewing coupons on cart-to-view ratios in the post-treatment period, estimated via the IV approach. Relative to the control condition, viewing shopping-cart promotions increased cart-to-view ratio by 0.52% for all sellers and 0.53% for no-promotion sellers (both p-values < 0.0001).

Altogether, we find that being exposed to the shopping-cart promotion program drove customers to strategically add products to the cart to obtain the same type of promotion again, and the average treatment effect of viewing coupons was stronger. We also observe that such effects spill over to no-promotion sellers, possibly because customers may infer that they can obtain the same type of promotion from sellers who had not targeted them before.

Table 4 Long-term Effects of Promotions on Strategic Customer Behavior

	Cart-to-view Ratio	Price Paid	Price Paid	Cart-to-view Ratio	Price Paid	Price Paid
Panel A: Long-term Intent-to-treat Effects	(All Sellers)			(No-promotion Sellers)		
	(1)	(2)	(3)	(4)	(5)	(6)
Treatment	0.0003****	-0.0075**	-0.0075**	0.0003****	-0.0065**	-0.0065**
	(0.0001)	(0.0033)	(0.0033)	(0.0001)	(0.0032)	(0.0032)
Relative Effect Size	0.41%	-0.02%	-0.02%	0.41%	-0.02%	-0.02%
Panel B: Long-term Average Treatment Effects of Viewing Coupons	Cart-to-view Ratio	Price Paid	Price Paid	Cart-to-view Ratio	Price Paid	Price Paid
	(All Sellers)			(No-promotion Sellers)		
	(1)	(2)	(3)	(4)	(5)	(6)
Ever Viewing Coupon	0.0004****	-0.0105**	-0.0105**	0.0004****	-0.0092**	-0.0091**
	(0.0001)	(0.0046)	(0.0046)	(0.0001)	(0.0045)	(0.0045)
Relative Effect Size	0.52%	-0.03%	-0.03%	0.53%	-0.02%	-0.02%
Date Fixed Effects	Yes	No	Yes	Yes	No	Yes
Product Fixed Effects	No	Yes	Yes	No	Yes	Yes
Observations	15,059,021	1,649,193	1,649,193	15,048,820	1,619,995	1,619,995

Note: $*p < 0.10; **p < 0.05; ***p < 0.01; ****p < 0.001$. Robust standard errors are in parentheses.

Note: This table shows the long-term effects of being in the treatment group (Panel A) and viewing shopping-cart coupons (Panel B) on cart-to-view ratio (Columns (1) and (3)) and price paid by customers (Columns (2) and (4)). Columns (1)-(3) include *all* sellers, and Columns (4)-(6) focus on *no-promotion* sellers. Columns (2) and (5) do not include date fixed effects, and Columns (3) and (6) control for date fixed effects.

6.2.2. Cross-mechanism Strategic Effect Next, we compare the price that treated and control customers paid for the same product. For this analysis, we have two data selection criteria. First, we select products that were purchased by at least 50 unique customers in the treatment condition and 50 unique customers in the control condition during our post-treatment period. This allows us to have enough observations of treated and control customers to compare prices that they paid for the same product while accurately controlling for product fixed effects. Our results are qualitatively similar if we use $10, 15, 30$, or 60 unique customers in each condition. Second, we focus on products that did not offer shopping-cart promotions in our post-treatment period because we are interested in testing whether customers in the treatment condition managed to obtain a cheaper price for the same product in other ways beyond getting a shopping-cart promotion. In the end, there were $4,794$ products satisfying these selection criteria, and these products were purchased $1,649,193$ times by $596,369$ customers in our sample.

Our unit of observation is one transaction ijk that customer i made for product j at transaction k. For our intent-to-treat analysis, we use the following regression specification to compare price paid between the treatment and control customers in the post-treatment period:

$$Price\ Paid_{ijk} = \theta_0 + \theta_1\, Treatment_i + X_j + D_t + \epsilon_{ijk} \tag{7}$$

where $Price\ Paid_{ijk}$ equals customer i's average price per item on transaction ijk. X_j represents product fixed effects. D_t represents date fixed effects. θ_1 is the coefficient of interest, which captures the average effect of being in the treatment group on the price paid by a customer for a product.

We start by analyzing all products that met our selection criteria. Column 2 in Table 4's Panel A reports the results from Specification (7) without controlling for date fixed effects. We find that for the same product, the price paid by treated customers is 0.02% lower than price paid by control customers on average (p-value < 0.05). Column 3 of Table 4's Panel A reports the results from specification 7, controlling for date fixed effects. Again, we find that for the same product, the price paid by treated customers is 0.02% lower than the price paid by control customers in the post-treatment period (p-value < 0.05). Furthermore, Columns 2 and 3 of Table 4's Panel B report the causal effect of viewing coupons on price paid for a product in the post-treatment period, which we estimate via the IV approach described in Section 5.2. Column 2 shows that, relative to the control condition, viewing shopping-cart promotions decreased price paid for a product by 0.03%, and Column 3 confirms that the results remain virtually unchanged when we control for date fixed effects (both p-values < 0.05).

Next, we turn to products from no-promotion sellers.[12] Column 5 of Table 4 shows that, relative to the control condition, exposure to our promotion program reduced the price customers paid for products from no-promotion sellers by 0.02% (Panel A) and viewing coupons decreased price paid by 0.02% (Panel B). The effects remain basically the same when we control for date fixed effects, as Column 6 shows.

These results provide evidence that treated customers were able to pay a lower price for the same product than control customers, via other promotion mechanisms beyond getting a shopping-cart promotion. Furthermore, we observe a similar magnitude of effect towards all sellers versus no-promotion sellers. As explained in Section 3.3, such a cross-mechanism strategic effect may arise on the Alibaba platform via two main strategies: (1) treated customers may be more likely to buy a product during a retailer-specific promotional event and (2) treated customers may be more likely to buy a product on a platform-wide shopping holiday. The estimated effects in Column 2 and 5 of Table 4 can be attributed to both strategies. In contrast, by including date fixed effects, Column 3 and 6 control for differences in dates when customers made purchases. Thus, Column 3 and 6 estimate the difference in price paid between treated and control customers while teasing out the influence of platform-wide shopping holidays. The fact that the coefficients are very similar with or without date fixed effects suggests that the majority of the decrease in price paid came from retailer-specific promotions. This is consistent with our sample period because it does not include Alibaba's two major platform-wide holidays, namely June 18 and November 11.

[12] In our analysis above, we have excluded all products offering shopping-cart promotions, but the sample still contains products from sellers who previously offered shopping-cart promotions for other products. For our analysis about no-promotion sellers, we further exclude products from sellers who previously offered shopping-cart promotions.

7. Heterogeneous Treatment Effects of Promotions

In this section, we report exploratory analyses about the heterogeneous treatment effects of our promotion program.

7.1. Heterogeneous Treatment Effects Across Promotion, Seller, and Consumer Characteristics

Based on past literature and our data availability, we identify a number of coupon, promotion seller, and consumer characteristics that may influence the short-term effectiveness of price promotions and the long-term effects of price promotions on consumer behaviors. The operationalization of these moderators can be found in Panel C of Table 7.

1. *Non-main-industry Promotion*: Prior research suggests that people are more likely to purchase sales items in response to unexpected coupons than expected coupons (Heilman et al. 2002). In our case, consumers may find it more surprising for sellers to offer discounts on products that do not belong to their main category. We explore whether consumers are more interested in seizing a promotion opportunity when the shopping-cart promotion is offered for a product not from the corresponding seller's main industry. In our long-term analysis, we explore whether receiving a shopping-cart promotion for a product not from the corresponding seller's main industry in treatment period is associated with greater changes in subsequent consumer behavior in the post-treatment period.

2. *Promotion Depth*: Since a larger discount indicates more cost saving and leaves consumers with a higher utility, a larger discount should be more likely to attract consumers. Also, past research suggests that consumers are more likely to update their expectations about future deals and prices after receiving a larger price cut than after receiving a smaller price cut (Kalwani and Yim 1992, Anderson and Simester 2004). Thus, we explore whether deeper promotions not only are more effective in lifting the short-term sales of promoted products but also have a greater long-term impact on consumer engagement and strategic behavior.

3. *Promotion Seller GMV*: Past research suggests that leading sellers that dominate the marketplace in its sales have higher market power and are more likely to affect customer behavior (Nevo 2001). Thus, we test whether promotions offered by sellers with a larger average daily GMV are more likely to lift the sales of promoted products and affect consumers' long-term behavior.

4. *Consumer Credit Rating*: On Alibaba, a customer's credit rating equals the number of positive reviews minus the number of negative reviews she has received; almost all reviews given to customers are positive. Thus, consumer credit rating is very similar to the number of transactions a consumer has made on the platform, reflecting consumer experience. Past research suggests that more experienced consumers are better at searching for product information and finding deals (Johnson and Russo 1984). In our case, more experienced customers on Alibaba may be more likely

to notice the shopping-cart promotions and thus respond to these promotions more strongly both in the short and long term. Therefore, we test whether the short- and long-term effects of our promotion program increase with customers' credit rating.

5. *Number of Promotions (i.e., Promotion Frequency)*: Prior research suggests that promotion frequency increases consumers' tendency to adjust their expectations about prices and future promotions (Kalwani and Yim 1992). We explore whether receiving more promotions during the treatment period, which reflects higher promotion frequency, is associated with a greater long-term impact on consumer engagement and strategy behavior in the post-treatment period.[13]

We rely on the following OLS regression specifications to test the interaction between the aforementioned list of factors and our promotion treatment in predicting short-term sales as well as customer engagement and strategic behaviors in the long term:

$$\begin{aligned}\text{Outcome Variable}_{ijt} =& \alpha_0 + \alpha_1 \text{Treatment}_i + \alpha_2 \text{Moderating Factor}_{ijt} + \\ & \alpha_3 \text{Treatment}_i * \text{Moderating Factor}_{ijt} + X_j + D_t + \epsilon_{ijt} \qquad (8)\\ \text{Outcome Variable}_{it} =& \alpha_0 + \alpha_1 \text{Treatment}_i + \alpha_2 \text{Moderating Factor}_i + \\ & \alpha_3 \text{Treatment}_i * \text{Moderating Factor}_i + D_t + \epsilon_{it} \qquad (9)\\ \textit{Price Paid}_{ijk} =& \alpha_0 + \alpha_1 \text{Treatment}_i + \alpha_2 \text{Moderating Factor}_i + \\ & \alpha_3 \text{Treatment}_i * \text{Moderating Factor}_i + X_j + D_t + \epsilon_{ijk} \qquad (10)\end{aligned}$$

where Moderating Factor_{ijt} and Moderating Factor_i are from Panel C of Table 7, and other variables in Specifications (8), (9), and (10) have the same definition as they do in Specifications (1)-(2), (3)-(6) and (7), respectively. Moderators that are continuous variables are standardized before entering into the regressions.

Table 5 reports the results from Specifications (8)-(10). In all models, the interaction term between treatment and a moderator is the variable of interest. Panel A of Table 5 shows that a promotion that did not belong to the corresponding seller's main industry (vs. a promotion from the corresponding seller's main industry) amplified the short-term sales lift brought by promotions (Column 1) but not the expenditure on promoted products (Column 2), intensified customers' long-term engagement (Columns 3-4) and direct strategic behavior (Column 5), and directionally but not statistically significantly decreased the price consumers paid in the long term (Column 6). Panel B indicates that larger promotions amplify the short-term sale lifts (Column 1) but abate the short-term expenditure lifts (Column 2). These results suggest that larger promotions attract

[13] This variable cannot be used as a moderator in our short-term analysis because it is inappropriate to test whether a consumer's response to a given promotion is affected by the total number of promotions that would target her during the whole treatment period.

Table 5 Heterogeneous Short- and Long-term Effects Across Promotions, Promotion Sellers, and Consumers

Panel A: Non-main-industry Promotion						
	P.I.	Expenditure	Daily (Views	P.I)	Cart Ratio	Price Paid
	(1)	(2)	(3)	(4)	(5)	(6)
Treatment	0.0115****	0.9371****	0.1483****	0.0006****	0.0003****	-0.0070**
	(0.0002)	(0.0259)	(0.0186)	(0.0002)	(0.0001)	(0.0033)
Treatment ×	0.0042***	0.0452	0.2721**	0.0032***	0.0013***	-0.0194
Non-main-industry Promotion	(0.0016)	(0.2012)	(0.1344)	(0.0012)	(0.0005)	(0.0224)
Panel B: Promotion Depth						
	P.I.	Expenditure	Daily (Views	P.I)	Cart Ratio	Price Paid
	(1)	(2)	(3)	(4)	(5)	(6)
Treatment	0.0115****	0.9474****	0.1830****	0.0008****	0.0003****	-0.0067**
	(0.0002)	(0.0268)	(0.0191)	(0.0002)	(0.0001)	(0.0033)
Treatment ×	0.0021****	-0.1518****	0.0719****	0.0003**	0.0001	-0.0050
Promotion Depth	(0.0002)	(0.0257)	(0.0190)	(0.0002)	(0.0001)	(0.0032)
Panel C: Promotion Seller GMV						
	P.I.	Expenditure	Daily (Views	P.I)	Cart Ratio	Price Paid
	(1)	(2)	(3)	(4)	(5)	(6)
Treatment	0.0116****	0.9435****	0.1817****	0.0008****	0.0003****	-0.0072**
	(0.0002)	(0.0268)	(0.0191)	(0.0002)	(0.0001)	(0.0033)
Treatment ×	0.0005***	0.1524****	0.0454**	0.0006****	0.0001	-0.0079**
Promotion Seller GMV	(0.0002)	(0.0260)	(0.0188)	(0.0002)	(0.0001)	(0.0033)
Panel D: Consumer Credit Rating						
	P.I.	Expenditure	Daily (Views	P.I)	Cart Ratio	Price Paid
	(1)	(2)	(3)	(4)	(5)	(6)
Treatment	0.0115****	0.9388****	0.1842****	0.0008****	0.0003****	-0.0060*
	(0.0002)	(0.0268)	(0.0189)	(0.0002)	(0.0001)	(0.0033)
Treatment ×	0.0044****	0.5634****	0.1764****	0.0023****	0.0003****	-0.0077****
Consumer Credit Rating	(0.0002)	(0.0266)	(0.0200)	(0.0002)	(0.0001)	(0.0019)
Panel E: Number of Promotions						
	P.I.	Expenditure	Daily (Views	P.I)	Cart Ratio	Price Paid
	(1)	(2)	(3)	(4)	(5)	(6)
Treatment			0.7654****	0.0024****	0.0034****	-0.0091**
			(0.0213)	(0.0002)	(0.0001)	(0.0036)
Treatment ×			2.0821****	0.0082****	0.0087****	-0.0040
Number of Promotions			(0.0300)	(0.0003)	(0.0001)	(0.0053)
Product Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Date Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1,613,049	1,613,049	24,010,020	24,010,020	15,059,021	1,649,193

Note: $*p < 0.10; **p < 0.05; ***p < 0.01; ****p < 0.001$. Robust standard errors are in parentheses.

Note: This table shows the heterogeneous treatment effects of being in the treatment group on purchase incidence of and expenditure on promoted products in the short term (Columns (1)-(2), respectively), daily number of products viewed (Column (3)), daily purchase incidence (Column (4)), cart-to-view ratio (Column (5)), as well as price paid for the same product (Column (6)) in the post-treatment period. For all columns, we only include observations that have known values across all moderators studied. Regression coefficients on moderating factors are not reported in this table due to space constraint.

more consumers to buy, but the additional increase in purchase incidence does not make up for the increased cost of promotion (at least in a linear specification). Also, we find that larger (vs. smaller) promotions increased customers' long-term engagement even more (Columns 3 and 4), and directionally intensified customers' long-term strategic behaviors (Columns 5 and 6).

Panel C shows that promotions offered by larger sellers (in terms of sellers' average daily GMV) more effectively increased the purchase incidence and revenue of promoted products than promotions offered by smaller sellers (Columns 1 and 2). Furthermore, promotions from sellers with a larger (vs. smaller) GMV increased daily views and purchase incidence, and decreased price paid for the same product in the post-treatment period to a greater extent (Columns 3, 4, and 6). The size of promotion sellers directionally but not statistically significantly amplified the effect of our treatment on cart-to-view ratio (Column 5).

Panel D shows that consumers with higher credit ratings (i.e., more experienced consumers) were more likely to increase their short-term purchase incidence and expenditure upon receiving a promotion (Columns 1-2) and increased their engagement in the long term to a greater extent after receiving a promotion (Columns 3-4). However, experienced (vs. novice) consumers also became more strategic and paid a lower price for the same product in the long term due to receiving a promotion (Columns 5 and 6). Panel E shows that consumers receiving more (vs. fewer) promotions in the treatment period increased their long-term engagement and direct strategic behavior to a greater extent (Columns 3-5), and directionally but not statistically significantly lowered the prices paid in the long term (Column 6) as a result of our promotion treatment.

Overall, we find that these moderating factors can affect our short-term or long-term treatment effects, which sheds light on how to better design promotion programs. The heterogeneous treatment effects of promotions on consumer behavior with respect to no-promotion sellers are qualitatively similar to those reported in Table 5.

7.2. Market Heterogeneous Effects Across All Sellers

In this section, we examine which type of seller is more likely to be affected by consumers' exposure to price promotions in the long term. Specifically, we classify all sellers on the platform into four groups of equal size based on their average daily GMV during the 30 days prior to the start of our study period. To make the size of our dataset manageable, we obtain weekly observations for each consumer. Specifically, for each consumer each week in the post-treatment period, we obtain the number of products she viewed (*Weekly Views*), the number of days on which she made any purchases (*Weekly Shopping Days*), her spending on the platform (*Weekly Expenditure*), and the proportion of products she added to her cart upon viewing them (*Weekly Cart-to-view Ratio*), with respect to each of the four groups of sellers. We use OLS regressions similar to Specifications (3) - (6) to separately estimate the impact of our promotion treatment on consumers' weekly views, shopping days, expenditure, and cart-to-view ratio in the post-treatment period with respect to

Table 6 Spillover Effects to Sellers of Different Sizes

	Weekly (Views	Shopping Days	Expenditure	Cart-to-view Ratio)
Panel A: Intent-to-treat Effects - Sellers in the Bottom Quartile of Daily GMV				
	(1)	(2)	(3)	(4)
Treatment	0.0132*	1e-05	-0.0886	0.0002
	(0.0075)	(0.0002)	(0.0848)	(0.0003)
Observations	3,724,016	3,724,016	3,724,016	1,720,651
Panel B: Intent-to-treat Effects - Sellers in the Second Quartile of Daily GMV				
Treatment	0.0261**	0.0002	0.0659	-2e-05
	(0.0120)	(0.0003)	(0.0681)	(0.0002)
Observations	3,724,016	3,724,016	3,724,016	2,184,238
Panel C: Intent-to-treat Effects - Sellers in the Third Quartile of Daily GMV				
Treatment	0.0636**	-0.0001	-0.2019	0.0004***
	(0.0292)	(0.0005)	(0.1250)	(0.0002)
Observations	3,724,016	3,724,016	3,724,016	2,831,574
Panel D: Intent-to-treat Effects - Sellers in the Top Quartile of Daily GMV				
Treatment	0.7690****	0.0036***	0.0210	0.0004****
	(0.1750)	(0.0013)	(0.5813)	(0.0001)
Observations	3,724,016	3,724,016	3,724,016	3,319,307

Note: $*p < 0.10; **p < 0.05; ***p < 0.01; ****p < 0.001$. Robust standard errors are in parentheses.

Note: This table shows the effects of being in the treatment group on the number of products viewed per week (Column (1)), the number of shopping days per week (Column (2)), weekly expenditure (Column (3)), and cart-to-view ratio per week (Column (4)) in the post-treatment period. This table separately reports regression results for four groups of sellers based on their GMV. Column (4) has fewer observations than Columns (1)-(3) because observations where a given consumer did not view any product from a given group of sellers in a given week were excluded from the regression reported in Column (4). All regressions include week fixed effects.

each of the four groups of sellers. For each group of sellers, we have a balanced panel including $3,724,016$ observations ($931,004$ customers $\times$ 4 weeks).[14] Results are reported in Table 6.

Column 1 of Table 6 shows that our promotion treatment increased product views among sellers in the bottom quartile (by 0.43%, p-value $= 0.08$), second quartile (by 0.44%, p-value < 0.05), third quartile (by 0.35%, p-value < 0.05), and top quartile (by 0.58%, p-value < 0.001), relative to the control condition. Column 2 shows that our promotion treatment only significantly increased the number of shopping days among sellers in the top quartile by 0.33% (p-value < 0.01) but did not significantly affect the number of shopping days among sellers with smaller values of GMV. Column 3 shows that our promotion treatment did not have a statistically significant impact on expenditure for sellers in any quartile. Lastly, Column 4 indicates that our promotion treatment significantly increased cart-to-view ratio among sellers in the top quartile (by 0.47%, p-value < 0.001) and third quartile (by 0.56%, p-value < 0.01), but did not significantly change cart-to-view ratio among sellers in the bottom and second quartiles. To summarize, in terms of consumer engagement as measured by product views and purchase frequency, sellers with the largest GMV seem to have been affected

[14] Five of the consumers in our original one-million consumers sample did not have data in this weekly-observation dataset. Thus, this analysis is based on 931,004 consumers, rather than 931,009 consumers as in our analysis about daily observations.

by our promotion treatment to the largest extent; in terms of consumers' direct strategic behavior, only sellers in the top and third quartiles were influenced by our promotion treatment.

8. Discussion and Conclusion

8.1. General Discussion

We conducted and analyzed a large field experiment that randomized whether customers were exposed to shopping-cart promotions on the world's largest retailing platform. We present several main findings. First, we show that our shopping-cart promotion program, which offers promotions to eligible customers for selected products in their shopping cart, is an effective way of personalizing prices. Upon receiving a coupon, customers, on average, increased their purchase likelihood of and expenditures on promoted products by 116% and 90%, respectively. In addition, viewing such shopping-cart coupons increased customers' purchase likelihood and expenditure by 173% and 135%, respectively.

Second, and more importantly, we demonstrate that the promotion program has unintended but positive consequences on customer engagement with the platform even after initial promotions have long expired. Specifically, being exposed to shopping-cart promotions increased the number of products customers viewed per day on the platform by 0.53% and the likelihood of purchasing any product on a given day by 0.29% during the month following our treatment period. In addition, viewing shopping-cart coupons in the treatment period subsequently increased the number of daily product views by 0.68% and purchase incidence by 0.37%.

Third, we find that customers exposed to price promotions became more strategic than control customers in two ways. The first manifestation is the direct strategic effect: being exposed to shopping-cart promotions and viewing coupons in the treatment period subsequently increased the ratio of products added to the cart conditional on views by 0.41% and 0.52%, respectively. The second manifestation is the cross-mechanism strategic effect: being exposed to shopping-cart promotions and viewing coupons in the treatment period subsequently decreased the price paid for a product that did *not* offer shopping-cart promotions by 0.02% and 0.03%, respectively.

Fourth, we observe that all of the long-term effects of promotions spilled over to a large set of no-promotion sellers on the platform. We find that the long-term effects of receiving and viewing shopping-cart promotions on customers' search and purchasing activity spilled over to retailers that did not previously offer shopping-cart promotions; the increased level of strategic behavior in the post-treatment period held true for customers' interaction with no-promotion sellers.

In addition, our exploratory analyses about heterogeneous treatment effects suggest that our promotion program generally affected consumer behavior to a greater extent in the short or long term when (1) promoted products did not belong to promotion sellers' main industry, (2) promotion depth was higher, (3) promotion sellers had a larger GMV, (4) consumers were more experienced, and (5) customer received promotions more frequently. We also find suggestive evidence that the long-term impact of our promotion program was stronger among sellers with a large GMV.

It is worth noting that our shopping-cart promotion program did not often give out large discounts (the average and median discount rates were 17% and 13%, respectively) and that our field experiment was not designed to focus on detecting the long-term effects of the promotion program. Rather, the size of a discount had to comply with Alibaba's and retailers' constraints, and the main purpose of the experiment from Alibaba's perspective was to understand how shopping-cart promotions influence sales in the short-term. This program is indeed effective in the short term: during our treatment period, we estimate that our observed 90% increase in expenditure could cause a total increase in spending on promoted products by 156 million RMB across 100 million consumers involved in the program if everyone had been treated. Even though the long-term impacts of promotions were not expected by Alibaba, we observe economically significant long-term unintended consequences as a result of those promotions. For example, during the one-month-long post-treatment period, we estimate that our promotion program could cause a total lift of 352 million product views and 162 million shopping days due to increased consumer engagement, induce consumers to add 20 million more products to their shopping cart due to direct strategic behavior alone, and lead consumers to pay 16 million RMB less due to cross-mechanism strategic behavior if all 100 million customers had been treated. See Online Appendix Section 4 for detail about our estimates.

We conducted further analyses to explore how our observed effects changed over the course of the post-treatment period as well as whether these effects lasted beyond the first month following our treatment period (See Online Appendix Section 3 for detailed methods and results). To address the first question, we separately ran regression specifications (3), (5), and (7) for each week during our one-month post-treatment period. We found that the effects of promotions on consumer engagement and strategic behavior were not statistically significantly distinguishable across weeks and that the effects did not wear off as the post-treatment month went on. We also obtained weekly observations about consumer behavior in the sixth month (September 12, 2016-October 9, 2016) and twelfth month (February 12, 2017-March 11, 2017) following our treatment period. We found that the positive effects of our shopping-cart promotion program on searches and purchase frequency persisted into the sixth month following our promotion treatment, but the strategic effects did not. Twelve months later, the effects of our promotion program on consumer engagement and strategic behavior wore off. Since most consumers (80%) were only targeted *once* by shopping-cart promotions in our treatment period, it is not surprising that not all effects of our promotion program on consumer behavior lasted beyond six months. Nonetheless, as our previous back-of-envelope calculation suggests, the overall impact of our promotion program on consumer behavior is economically meaningful even when we only consider the first month after our promotion treatment.

8.2. Practical Implications

Our research offers several important managerial implications. First, our findings highlight that platforms and retailers should be very cautious about the negative effects of implementing price promotion programs on customers' long-term behaviors. The price promotion program we studied not only drove customers to add a greater proportion of viewed products to the cart in anticipation of the same type of price promotions they enjoyed in the past, but also led them to pay less for the same products in general. Understanding how customers develop strategic behavior with respect to different pricing policies is critical for retailers and platform managers. For example, they should not only consider the immediate cost of running a promotion on revenue, but also understand the long-term costs of promotions due to customers becoming more strategic.

Second, although we document that our price promotion program intensifies strategic customer behavior, it also has an important positive long-term consequence: it increases customers' engagement with the platform. Customers' willingness to spend time on a platform is critical to platform success. Therefore, when retailers and platform managers implement dynamic pricing programs through promotions, they should consider how to take advantage of customers' increased interaction with the platform. For example, as customers increase the time they spend searching for "good deals" on a platform, the platform can design better personalized recommendations and turn the increased search traffic into purchasing.

Third, our findings suggest that retailing platform managers should be cautious about the spillover effects of dynamic pricing programs. We show that seemingly small promotions from some retailers can create a non-trivial spillover effect to other retailers on the platform. Therefore, when designing dynamic pricing policies on platforms, managers should consider how such policies may affect the retailers who will adopt the policies as well as other seemingly unrelated sellers.

In addition, our heterogeneous treatment effects suggest that platform managers and retailers may design better promotion programs by being thoughtful about what type of promotion to offer, what type of consumer to target, and what type of seller to enroll in a promotion program. For example, if platform managers hope to maximize short-term profits of promoted products, our results suggest that they may be better off by giving larger retailers access to a promotional tool or targeting more experienced consumers. If platform managers primarily care about long-term consumer behavior on the entire platform, our findings suggest that offering surprising promotions or deeper promotions could amplify the positive effects of promotions on consumer engagement without decreasing the price consumers paid in the long term. Also, platforms face tradeoffs (a) between short-term and long-term impacts of promotions and (b) between positive and negative long-term effects when they choose parameters and scopes for a promotion program. For example, offering deeper promotions diminishes the short-term revenue lifts brought by promotions, but amplifies the positive effects of promotions on consumer engagement; selecting larger sellers or targeting more experienced consumers can further boost the benefits of promotions on consumer engagement, while intensifying the negative strategic effect. The right choice for a platform depends on its objective function (e.g., the weight on short-term revenue vs. long-term revenue stream) and the extent to which increased consumer engagement can be successfully monetized.

Implications for Alibaba. After our experiment and analysis, Alibaba's senior management group became convinced that dynamic pricing through offering price promotions could make customers more strategic on the entire platform. Therefore, they become more cautious about launching other targeted promotion programs after our experiment. They also decided not to aggressively expand this shopping-cart program to more retailers despite its impressive short-term benefits on sales. Moreover, the senior management group was happy to see that promotions can improve customer engagement on the platform in the long term, and found it helpful to understand how consumers' reactions to price promotions could vary based on characteristics of promotions offered, consumers targeted, or retailers participating in a promotion program. Combining all these considerations, the management group, at our suggestion, is trying to use non-monetary incentives to help retailers converge customers who have already added products to carts without purchasing them for some time. In addition, motivated by our results, Alibaba has started designing personalized promotion programs using deep reinforcement learning to maximize the positive short-term lifts while minimizing the negative long-term strategic behaviors.

8.3. Limitations and Future Research

We studied the effects of offering targeted price promotions based on customer and product characteristics, a very popular way of implementing dynamic pricing. There are also other dynamic pricing programs that directly manipulate prices instead of giving out promotions (such as the programs examined in Xu et al. 2016, Cheung et al. 2017, Fisher et al. 2017) or price discriminate not only based on other factors beyond customer and product characteristics (e.g., purchasing time; Özer et al. 2012). While most of all customers observe one price change by receiving one promotion in our program, some other dynamic pricing programs based on multi-arm bandit algorithms may change prices much more frequently to identify the optimal price (e.g., Cheung et al. 2017). Therefore, it is an interesting question whether our findings can generalize to other types of dynamic pricing programs.

Moreover, though the gender and age of consumers targeted by shopping-cart promotions were comparable to those of general consumers using Alibaba's mobile app (see section 4.1), it is plausible that customers targeted by our promotion program had high interest in the promoted products to begin with, given their initial decision to add them to their carts. Customers with low interest in promoted products might find price promotions less attractive and salient; as a result, these customers may not become more strategic or engaged with the platform after receiving promotions. Therefore, an interesting avenue for future research is to investigate how different types of customers, besides the consumer characteristics that we examined in this paper, react to different dynamic pricing policies and promotion strategies in the long term. In particular, considering that price sensitivity influences how consumers respond to market factors (e.g., Gaur and Fisher 2005, Lu et al. 2013, Buell et al. 2016), it is useful for future research to explore how to personalize promotions for customers with different levels of price sensitivity in order to maximize the positive effects while minimizing the negative effects of promotions on customers' long-term behaviors.

One important observation from our study is that customers may become more strategic in the long term after being exposed to a promotion program. It is important to incorporate this observation into future theoretical models and understand how to optimize operations decisions, not only when customers are strategic but also when customers may learn to become more strategic over time. A second important observation from our study is that the effects of pricing policies from certain sellers will have substantial spillover effects to other sellers on the same platform. An interesting direction for future research would be to investigate how such spillover effects may affect the incentive alignment among sellers and the platform and, more importantly, how platforms can incorporate such spillover effects in their operations decision making, which may have implications for price dispersion and market efficiency (Parker et al. 2016).

References

Ailawadi, Kusum L, Scott A Neslin. 1998. The effect of promotion on consumption: Buying more and consuming it faster. *Journal of Marketing Research* 390–398.

Allon, Gad, Dennis J Zhang. 2015. Managing service systems in the presence of social networks, *working paper.*

Anderson, Eric T, Duncan I Simester. 2004. Long-run effects of promotion depth on new versus established customers: Three field studies. *Marketing Science* **23**(1) 4–20.

Angrist, Joshua D, Jörn-Steffen Pischke. 2008. *Mostly harmless econometrics: An empiricist's companion.* Princeton university press.

Broder, Josef. 2011. Online algorithms for revenue management, *phd thesis.*

Buell, Ryan W, Dennis Campbell, Frances X Frei. 2016. How do customers respond to increased service quality competition? *Manufacturing & Service Operations Management* **18**(4) 585–607.

Burdett, Kenneth, Kenneth L Judd. 1983. Equilibrium price dispersion. *Econometrica: Journal of the Econometric Society* 955–969.

Cachon, Gérard P, Pnina Feldman. 2015. Price commitments with strategic consumers: Why it can be optimal to discount more frequently than optimal. *Manufacturing & Service Operations Management* **17**(3) 399–410.

Cheung, Wang Chi, David Simchi-Levi, He Wang. 2017. Dynamic pricing and demand learning with limited price experimentation. *Operations Research* .

Cohen, Maxime C, Ngai-Hang Zachary Leung, Kiran Panchamgam, Georgia Perakis, Anthony Smith. 2017. The impact of linear optimization on promotion planning. *Operations Research* **65**(2) 446–468.

Cui, Ruomeng, Jun Li, Dennis J Zhang. 2016. Discrimination with incomplete information in the sharing economy: Field evidence from airbnb, *working paper*.

Ferreira, Kris Johnson, Bin Hong Alex Lee, David Simchi-Levi. 2015. Analytics for an online retailer: Demand forecasting and price optimization. *Manufacturing & Service Operations Management* **18**(1) 69–88.

Fisher, Marshall, Santiago Gallino, Jun Li. 2017. Competition-based dynamic pricing in online retailing: A methodology validated with field experiments. *Management Science* .

Fong, Nathan M, Yuchi Zhang, Xueming Luo, Xiaoyi Wang. 2016. Targeted promotions and cross-category spillover effects, *working paper*.

Gaur, Vishal, Marshall L Fisher. 2005. In-store experiments to determine the impact of price on sales. *Production and Operations Management* **14**(4) 377–387.

Heilman, Carrie M, Kent Nakamoto, Ambar G Rao. 2002. Pleasant surprises: Consumer response to unexpected in-store coupons. *Journal of Marketing Research* **39**(2) 242–252.

Jedidi, Kamel, Carl F Mela, Sunil Gupta. 1999. Managing advertising and promotion for long-run profitability. *Marketing science* **18**(1) 1–22.

Johnson, Eric J, J Edward Russo. 1984. Product familiarity and learning new information. *Journal of consumer research* **11**(1) 542–550.

Kabra, Ashish, Elena Belavina, Karan Girotra. 2015. Bike-share systems: Accessibility and availability. *Chicago Booth Research Paper* (15-04).

Kahneman, Daniel, Amos Tversky. 1979. Prospect theory: An analysis of decision under risk. *Econometrica: Journal of the econometric society* 263–291.

Kalwani, Manohar U, Chi Kin Yim. 1992. Consumer price and promotion expectations: An experimental study. *Journal of marketing Research* 90–100.

Kőszegi, Botond, Matthew Rabin. 2006. A model of reference-dependent preferences. *The Quarterly Journal of Economics* **121**(4) 1133–1165.

Lattin, James M, Randolph E Bucklin. 1989. Reference effects of price and promotion on brand choice behavior. *Journal of Marketing Research* **26**(3) 299–310.

Li, Jun, Nelson Granados, Serguei Netessine. 2014. Are consumers strategic? structural estimation from the air-travel industry. *Management Science* **60**(9) 2114–2137.

Li, Jun, Antonio Moreno, Dennis J Zhang. 2016. Pros vs joes: Agent pricing behavior in the sharing economy, *working paper*.

Lu, Yina, Andrés Musalem, Marcelo Olivares, Ariel Schilkrut. 2013. Measuring the effect of queues on customer purchases. *Management Science* **59**(8) 1743–1763.

Luo, Xueming, Michelle Andrews, Zheng Fang, Chee Wei Phang. 2014. Mobile targeting. *Management Science* **60**(7) 1738–1756.

Mela, Carl F, Sunil Gupta, Donald R Lehmann. 1997. The long-term impact of promotion and advertising on consumer brand choice. *Journal of Marketing research* 248–261.

Mela, Carl F, Kamel Jedidi, Douglas Bowman. 1998. The long-term impact of promotions on consumer stockpiling behavior. *Journal of Marketing research* 250–262.

Moon, Ken, Kostas Bimpikis, Haim Mendelson. 2017. Randomized markdowns and online monitoring. *Management Science* .

Nasiry, Javad, Ioana Popescu. 2011. Dynamic pricing with loss-averse consumers and peak-end anchoring. *Operations research* **59**(6) 1361–1368.

Natter, Martin, Thomas Reutterer, Andreas Mild, Alfred Taudes. 2007. Practice prize reportan assortmentwide decision-support system for dynamic pricing and promotion planning in diy retailing. *Marketing Science* **26**(4) 576–583.

Netessine, Serguei. 2006. Dynamic pricing of inventory/capacity with infrequent price changes. *European Journal of Operational Research* **174**(1) 553–580.

Nevo, Aviv. 2001. Measuring market power in the ready-to-eat cereal industry. *Econometrica* **69**(2) 307–342.

Nijs, Vincent R, Marnik G Dekimpe, Jan-Benedict EM Steenkamps, Dominique M Hanssens. 2001. The category-demand effects of price promotions. *Marketing science* **20**(1) 1–22.

Özer, Özalp, Ozalp Ozer, Robert Phillips. 2012. *The Oxford handbook of pricing management*. Oxford University Press.

Parker, Chris, Kamalini Ramdas, Nicos Savva. 2016. Is it enough? evidence from a natural experiment in india's agriculture markets. *Management Science* **62**(9) 2481–2503.

Pauwels, Koen, Dominique M Hanssens, Sivaramakrishnan Siddarth. 2002. The long-term effects of price promotions on category incidence, brand choice, and purchase quantity. *Journal of marketing research* **39**(4) 421–439.

Popescu, Ioana, Yaozhong Wu. 2007. Dynamic pricing strategies with reference effects. *Operations Research* **55**(3) 413–429.

Stigler, George J. 1961. The economics of information. *Journal of political economy* **69**(3) 213–225.

Su, Xuanming. 2007. Intertemporal pricing with strategic customer behavior. *Management Science* **53**(5) 726–741.

Su, Xuanming, Fuqiang Zhang. 2008. Strategic customer behavior, commitment, and supply chain performance. *Management Science* **54**(10) 1759–1773.

Tang, Christopher S, Jiaru Bai, Kut C So, Xiqun Michael Chen, Hai Wang. 2016. Coordinating supply and demand on an on-demand platform: Price, wage, and payout ratio .

Taylor, Terry. 2016. On-demand service platforms. *working paper* .

Thaler, Richard. 1985. Mental accounting and consumer choice. *Marketing Science* **4**(3) 199–214.

Wu, Jianghua, Ling Li, Li Da Xu. 2014. A randomized pricing decision support system in electronic commerce. *Decision Support Systems* **58** 43–52.

Xu, Joseph Jiaqi, Peter Fader, Senthil K Veeraraghavan. 2016. The revenue impact of dynamic pricing policies in major league baseball ticket sales, *working paper*.

Zhang, Dennis J, Gad Allon, Jan A Van Mieghem. 2017. Does social interaction improve learning outcomes? evidence from field experiments on massive open online courses. *Manufacturing & Service Operations Management* **19**(3) 347–367.

Appendices

Table 7 Summary of Key Dependent and Independent Variables Examined in This Study

Panel A: Key Dependent Variables

Dependent Variable	Definition	Section
Purchase Incidence$_{ijt}$ (P.I.)	A binary variable that equals 1 if customer i purchased promoted product j on day t, and 0 otherwise	Section 5
Expenditure$_{ijt}$	The total amount of money in RMB that customer i spent on promoted product j on day t	Section 5
Daily Views$_{it}$	The total number of products that customer i viewed on day t on the Alibaba platform	Section 6, 7
Daily Purchase Incidence$_{it}$ (Daily P.I.)	A binary variable that equals 1 if customer i made any purchases on day t on the Alibaba platform	Section 6, 7
Daily Expenditure$_{it}$	The total amount of money in RMB that customer i spent on the Alibaba platform on day t	Section 6, 7
Cart-to-view Ratio$_{it}$	The proportion of products that customer i added to her shopping cart on day t after viewing them	Section 6, 7
Price Paid$_{ijk}$	The average price customer i paid per item when she made a transaction k about product j	Section 6, 7

Panel B: Key Independent Variables

Independent Variable	Definition	Section
Treatment$_i$	A binary variable that equals 1 if customer i is in the treatment condition, and 0 otherwise	Section 5, 6, 7
Viewing Coupon$_{ijt}$	A binary variable that equals 1 if customer i viewed the shopping-cart coupon of promoted product j on day t, and 0 otherwise	Section 5
Ever Viewing Coupon$_i$	A binary variable that equals 1 if customer i at least viewed one coupon during the treatment period	Section 6

Panel C: Key Moderators

Moderator	Definition	Section
Non-main-industry Promotion$_{ijt}$	A binary variable that equals 1 if promotion j on day t was not from the corresponding seller's main industry	Section 7
Non-main-industry Promotion$_i$	A binary variable that equals 1 if the first promotion that targeted customer i was not from the corresponding seller's main industry	Section 7
Promotion Depth$_{ijt}$	A numerical variable between 0 and 1 that represents the depth of the promotion j for customer i on day t	Section 7
Promotion Depth$_i$	A numerical variable between 0 and 1 that represents the average promotion depth customer i experienced in the treatment period	Section 7
Promotion Seller GMV$_{ijt}$	A numerical variable representing the average daily gross merchandise value (GMV) of the promotion seller across the last 30 days prior to our study period	Section 7
Promotion Seller GMV$_i$	A numerical variable representing the average daily GMV of all promotion sellers that targeted customer i in the treatment period	Section 7
Customer Credit Rating$_i$	A numerical variable equaling the number of good reviews minus the number of bad reviews that customer i had received by the first day of our study period (i.e., February 12, 2016)	Section 7
Number of Promotions$_i$	A numerical variable representing how many times customer i received shopping-cart promotions in the treatment period	Section 7