

NBER WORKING PAPER SERIES

COMPETITIVE PRICE TARGETING WITH SMARTPHONE COUPONS

Jean-Pierre H. Dubé
Zheng Fang
Nathan Fong
Xueming Luo

Working Paper 22067
<http://www.nber.org/papers/w22067>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
March 2016

Dubé: Booth School of Business, University of Chicago, 5807 S Woodlawn Ave, Chicago, IL 60637 (e-mail: jdube@chicagobooth.edu); ; Fang: School of Business, Sichuan University (email: 149281891@qq.com); ; Fong: Fox School of Business, Temple University, Philadelphia, PA 19122 (e-mail: nmfong@temple.edu); Luo: Fox School of Business, Temple University, Philadelphia, PA 19122 (e-mail: luoxm@temple.edu). We would like to thank Eric Bradlow, Xiao Liu, Raji Srinivasan, K. Sudhir and Chunhua Wu for comments. We would also like to thank attendees of the 2015 Marketing Science Conference, MIT CODE Conference, and the NYU Big Data Conference, and workshop participants at Boston University, the Global Center for Big Data in Mobile Analytics, the University of Louisville and the University of California San Diego for their insightful comments. We are grateful for the cooperation of our mobile service provider with the implementation of our field experiment. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2016 by Jean-Pierre H. Dubé, Zheng Fang, Nathan Fong, and Xueming Luo. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Competitive Price Targeting with Smartphone Coupons
Jean-Pierre H. Dubé, Zheng Fang, Nathan Fong, and Xueming Luo
NBER Working Paper No. 22067
March 2016
JEL No. L1,L11,L13,M31

ABSTRACT

We conduct a large-scale field experiment to study competitive price discrimination in a duopoly market with two rival movie theaters. The firms use mobile targeting to offer different prices based on location and past consumer activity. A novel feature of our experiment is that we test a range of relative ticket prices from both firms to trace out their respective best-response functions and to assess equilibrium outcomes. We use our experimentally-generated data to estimate a demand model that can be used to predict the consumer choices and corresponding firm best-responses at price levels not included in the test. We find an empirically large return on investment when a single firm unilaterally targets its prices based on the geographic location or historical visit behavior of a mobile customer. However, these returns can be mitigated by competitive interactions whereby both firms simultaneously engage in targeting. In practice, firms typically test only their own prices and do not consider the competitive response of a rival. In our study of movie theaters, competition enhances the returns to behavioral targeting but reduces the returns to geo-targeting. Under geographic targeting, each theater offers a discount in the other rival's local market, toughening price competition. In contrast, under behavioral targeting, the strategic complementarity of prices coupled with the symmetric incentives of the two theaters to raise prices charged to high-recency customers softens price competition. Thus, managers need to consider how competition moderates the profitability of price targeting. Moreover, field experiments that hold the competitor's actions fixed may generate misleading conclusions if the permanent implementation of a tested action would likely elicit a competitive response.

Jean-Pierre H. Dubé
University of Chicago
Booth School of Business
5807 South Woodlawn Avenue
Chicago, IL 60637
and NBER
jdube@chicagobooth.edu

Zheng Fang
Sichuan University
29 Wangjiang Road
Chengdu, Sichuan
P.R.C 610064
149281891@qq.com

Nathan Fong
Temple University
Fox School of Business, Alter 517
1801 Liacouras Walk
Philadelphia, PA 19122
nmfong@temple.edu

Xueming Luo
Temple University
Fox School of Business, Room 511
1801 Liacouras Walk
Philadelphia, PA 19122
luoxm@temple.edu

1 Introduction

A large body of literature dating back at least to Pigou (1920) has studied the theory of monopoly price discrimination for a firm with market power (see Varian (1989) for a detailed overview). In general, so long as a firm has market power, consumers can be segmented, and arbitrage through resale is infeasible, a firm will typically have an incentive to price discriminate holding other firms' actions fixed. Focusing on targeted pricing to a group of customers, or "third-degree price discrimination", a monopoly firm's ability to target different prices to different consumers based on available consumer information weakly increases its profits. While it is tempting to apply the intuition from monopoly price discrimination to the analysis of oligopolistic markets, the intuition can be misleading (see Stole (2007) for a comprehensive discussion of oligopoly price discrimination). When oligopoly firms adopt targeted pricing strategies, the impact on industry profits depends on the gains from surplus extraction relative to any potential losses associated with the intensity of competition. Unlike the monopoly case, oligopoly price discrimination is more nuanced and the likely gains/losses to firms relative to uniform pricing depends on the characteristics of the market and the nature of price discrimination. Therefore, a firm empirically analyzing the likely returns to adopting price targeting technology in its product market could inadvertently overestimate the return on investment if it ignores its competitors' incentives to adopt similar technologies.

We empirically investigate the returns to targeted pricing in a competitive market through a study involving a large-scale, field experiment on mobile telephones in a large Asian city. We use a real-time subject pool of 18,000 mobile subscribers located within "geofences" centered on two shopping malls located 4 km (about 2.4 miles) apart, each with a competing movie theater chain¹. The experiment is conducted during an "off-peak" hour to avoid exceeding theater capacity. Each subject in our experiment receives an offer via SMS (short message service) from each of the two movie theaters, or receives a single SMS from one of the two movie theaters. Each SMS contains an offer to purchase a voucher to see a movie at that theater for a randomly chosen discount off the regular box-office price. A control group receives no offers and, hence, only has access to the regular ticket prices at the box office. A novel feature of this experimental design is that we randomize the prices charged by each of the two theaters to a given consumer. We can therefore trace out portions of each firm's best-response function². The best-response analysis allows us to compare each firm's pricing incentives when it unilaterally explores targeted pricing, as in monopoly price discrimination theory, versus when it considers competitive response, as in oligopoly price discrimination theory.

The mobile phone context enables several novel targeting opportunities. The ability to position a consumer geographically in real time using GPS and cell tower triangulation on mobile signal reception allows firms to implement geographic price discrimination. Assuming consumers incur travel costs and/or opportunity costs of time, one would expect a consumer's willingness-to-pay for a ticket to in-

¹A geofence is a digital perimeter defining the geographic boundaries of a market. In our setting, we use a 500 meter radius around each of the two shopping malls.

²The best-response function traces out the a given firm's optimal price for each of the possible prices charged by its rival.

crease with proximity to the theater. Similarly, the ability to track individual consumers through their unique telephone numbers allows firms to target prices based on a consumer's historic economic activity. For instance, the fact that a consumer has previously visited a movie theater likely reveals a higher willingness-to-pay for movie tickets relative to a consumer who has not recently visited a theater. Our experimental design enables us to test these location-based and activity-based theories of price discrimination.

As expected from the theory, the experiment shows that either firm would profit by offering deep discounts off their regular prices to consumers located near their competitor, which we refer to as "offensive" promotions. We use the term "offense" to indicate that the discount is geared towards poaching a customer close to the competitor's theater. However, the returns to such geographic targeting are mitigated if the competitor responds with targeted pricing to the same consumers, which we refer to as "defensive" promotions. We use the term "defense" to indicate that the discount is geared towards a theater's local customers, who represent its core business. We treat "competitive response" in the sense of a simultaneous game in this analysis. Interestingly, the response rate to defensive promotions appear to be relatively immune to the incidence and magnitude of the offensive promotions, suggesting an asymmetry in cross-promotional effects between the defensive and the offensive firm.

To analyze the mobile promotions as a non-cooperative strategic game, we use the results of the experiment to trace out portions of each firm's best-response function. Assuming each firm sets its price to maximize its current profit, the best-responses indicate how each firm would likely set its prices in response to the competitor's price decision. Using the discrete grid of price points tested in the experiment, our results suggest that both firms will offer deep discounts to all customer segments in equilibrium. Unfortunately, this result confounds the benefits of targeting and simple optimization. Even though we tested fairly large discounts off the firms' regular box-office prices, each firm's best responses lie below the range of tested prices. For this reason, we need a model to extrapolate the remainder of each firm's best-response function.

To solve this problem, we fit a demand model with flexible substitution patterns to the experimentally generated consumer choice data. A novel feature of our experiment is that it eliminates all the usual price endogeneity concerns that have challenged the traditional demand estimation literature using observational field data for consumer goods (e.g. Berry (1994)). To ensure the model can accommodate the observed asymmetric cross-price effects of defensive and offensive price promotions, we use a multinomial probit with correlated utility errors. We then use the estimated probit demand model to derive each firm's best-response function under various pricing scenarios.

The equilibrium uniform prices for both firms are about 70% lower than the regular box office prices. Recall that these prices apply to the off-peak time slot in contrast with the box office prices which are uniform across all time slots, including peak times. Turning to targeting, we first analyze the unilateral case whereby one firm targets prices while the other firm continues to charge its uniform price. Both firms unilaterally benefit from targeting on geographic location and/or on past consumer visit behavior, although the unilateral gains from geographic targeting are considerably larger than from behavioral

targeting. There is no theoretical reason for geographic targeting to dominate behavioral targeting per se. This finding is an empirical consequence of the differences in the degree of consumer heterogeneity in geographic space versus purchase recency.

As expected from the theory, the returns to targeting are quite different in equilibrium when both firms can endogenously choose whether or not to target. Interestingly, targeting on past behavior is more profitable in equilibrium than under unilateral targeting. This is due to the best-response symmetry in the consumer type segments and the fact that competition is more intense in the “High” market (consumers who recently visited the theater) than the “Low” market (consumers who did not recently visit the theater). In contrast, the profitability of location targeting falls in equilibrium relative to the unilateral cases. While both firms are still better off than under uniform pricing, the gains are mitigated by their asymmetric pricing incentives, leading each firm to target much lower prices in its rival’s local market. These results demonstrate how the unilateral manner in which many firms test targeting opportunities in practice could easily mis-estimate the benefits from targeting if there is competitive response. In our study, firms would over-estimate the returns to geo-targeting and would under-estimate the returns to behavioral targeting.

In spite of the fact that the returns to competitive price discrimination is largely an empirical question, there is a surprising lack of empirical research. A likely explanation is the practical difficulty of obtaining viable data. Since the returns to a firm’s targeting investments are moderated by the actions of its rival, it is critical to assess a firm’s best-responses. However, competing firms’ prices are rarely independently varied in such a way that allows econometric identification of competitive price interactions. In most markets, it would be illegal for competing firms to coordinate their pricing decisions even for the purposes of experimentation. In principle, one could run such an experiment in a retail setting, where the retailer could vary the prices of all the competing brands simultaneously. In this vein, our study takes advantage of a unique opportunity to observe the outcome of randomly assigned prices for two competing firms that operate completely independently of each other. A large, wireless service provider that operates a mobile marketing platform and issues promotions on behalf of its clients conducted the price experiment. In this regard, our experimental design points to opportunities to resolve the lack of evidence for the returns to price discrimination in competitive markets. More generally, our comparison of unilateral and equilibrium targeting demonstrates the importance and difficulty of evaluating the moderating effect of competitive response when optimizing marketing tactics. Doing so requires the combination of the right experiment and a model to relate the results to competitive theory. As a more general matter, our evidence confirms the fact that targetability need not lead to a profitable opportunity in a competitive market. The profitability of targeting is ultimately an empirical question about the nature of demand and competition.

2 Background

2.1 Competitive price targeting

Our empirical findings contribute to the theoretical literature on competitive third-degree price discrimination. For a monopolist, price discrimination will always weakly increase the firm's profits. Similarly, price discrimination typically weakly increases an oligopoly firm's profits, holding competitors' actions fixed. But, except under very stylized modeling assumptions, it is difficult to predict whether equilibrium profits rise or fall under competitive price discrimination. This difficulty is nicely demonstrated in Corts (1998), who makes an interesting distinction between two types of models. Suppose there are two consumer markets. A firm characterizes one of the markets as "weak" (and the other market as "strong") if, for any uniform price set by a competitor, the optimal price is lower than in the other market. A pricing model is characterized as exhibiting "best-response symmetry" if firms agree on the strong and weak markets. Otherwise, the model is characterized as exhibiting "best-response asymmetry." Under best-response symmetry, several papers have derived conditions under which the monopoly predictions appear to hold and price discrimination can increase profits under sufficiently intense competition in the "strong" market (Borenstein (1985); Holmes (1989); Armstrong and Vickers (2001)). Under best-response asymmetry, several stylized applications of the Hotelling model appear to predict an unambiguous prisoner's dilemma whereby all firms endogenously commit to price discriminating and generate lower equilibrium profits than under uniform pricing (Thisse and Vives (1988); Shaffer and Zhang (1995)). However, Corts (1998) shows that this result is not general and that, under best-response asymmetry, the uniform equilibrium prices need not lie between the price discrimination prices. In fact, best-response asymmetry turns out to be a necessary condition for two polar outcomes: "all out price competition" or "all out price increases." Under the former, prices and profits fall in all markets. Under the latter, prices and profits increase in all markets. Whether all out competition or all out price increases emerges is ultimately an empirical question regarding the relative importance each firm attaches to the strong and weak markets. Based on these findings, our approach to assessing the returns from targeting consists of devising a field experiment with which to assess the profitability of each consumer market to each of the competing firms.

More recently, Chen, Li, and Sun (2015) study the equilibrium incentives of firms to target prices by location, as opposed to by consumer. A novel feature of this setting is that consumers can endogenously move between locations based on their expectations about firms' geo-targeting incentives. This "cherry-picking" intensifies price competition so that, in equilibrium, a firm does not successfully poach its rival's local consumers³. We do not consider the ability of consumers to cherry pick in our empirical analysis. However, this would be an interesting topic for future research on geotargeting and consumers' strategic incentives.

³The authors also relax the usual "full market coverage" assumption by including a mass of marginal consumers who would not buy from either firm in the uniform price equilibrium. As long as the category expansion effects are not too large and price competition is not "too strong" in this neutral market, equilibrium profits can still increase under location targeting.

A related literature has analyzed the inter-temporal incentives for competing firms to target prices based on past consumer behavior (for a survey see Fudenberg and Villas-Boas (2006)). When consumers are also forward-looking, firms may find themselves in a prisoner's dilemma with lower profits than if they could credibly commit to not targeting based on past behavior. In our mobile campaigns, we do consider targeting based on past consumer visit behavior; but we do not consider the dynamic incentives of firms or consumers. Shin and Sudhir (2010) have found that even with forward-looking consumers, the prisoner's dilemma may not arise if consumers exhibit a sufficiently strong stochastic preference component like the one in our probit demand model.

Ex ante, our empirical setting appears to exhibit the intuitive properties of best-response asymmetry: the firms are geographically differentiated and can use mobile marketing to target consumers located close to their competitor. Calibrating a model of competitive pricing on our experimental results, we can observe whether the decision to adopt price targeting leads to a prisoner's dilemma. The presence of a prisoner's dilemma would empirically demonstrate how the presence of competition can reverse the profitability of price targeting. The lack of a prisoner's dilemma would not falsify the theory; however, it would suggest that competitive effects need to be quite severe in order for price targeting to lower profits, and would potentially demonstrate the insufficiency of best response asymmetry for generating such a result.

Several authors have conducted empirical tests for the incidence of competitive price discrimination (e.g. Shepard (1991); Borenstein and Shepard (1994); Goldberg and Verboven (2005); Busse and Rysman (2005)). Borenstein and Rose (1994) find that the degree of price discrimination in airline fares increases with the degree of competition. However, few papers have analyzed the profit implications of price discrimination and the potential, under best-response asymmetry, for all-out competition. In a study of the US RTE Cereal industry, Nevo and Wolfram (2002) find that shelf prices tend to be lower during periods of coupon availability. Besanko, Dubé, and Gupta (2003) conduct a structural analysis that calibrates a targeted couponing model with competing firms and manufacturers using Ketchup data, finding that competitive price targeting does not lead to all out war. However, their model also incorporates several other factors including a combination of horizontal and vertical differentiation between products, and horizontal and vertical competition between firms (retailers and manufacturers). Building on these findings, Pancras and Sudhir (2007) study the equilibrium incentives for a consumer data intermediary to sell access to customer data and provide targeting services to competing firms in a retail distribution channel. They too find that competitive targeting need not lead to all out war. Our setting provides a convenient context for studying competitive price discrimination as we have two firms selling relatively homogeneous products that are differentiated primarily along a single geographic dimension. We do not consider the incentives of the data intermediary, in this case the mobile platform, to sell targeting services. The platform offers targeting capabilities that use both real-time location, providing our horizontal dimension, and historical location, used to infer past behavior that comprises our vertical dimension.

2.2 Mobile marketing

Mobile technology has profoundly altered online consumer behavior and created new opportunities for targeted marketing. In particular, users of mobile devices tend to carry them at all times. Compared with PC-based internet access, a device is more likely to be tied to single user. Finally, the devices themselves offer location-specific services, and in many cases the service providers will receive location information. These features present an improved opportunity for targeting based on consumers' real-time locations, and on behavioral histories tied to a specific person. They also improve managers' ability to evaluate the effectiveness of marketing tactics, by providing improved measurement of individuals' behavior and the ability to run randomized experiments at the individual level.

Industry experts routinely report impressive response rates and incremental returns to firms that geo-target mobile offers. RocketFuel, a leading US provider of geotargeted mobile ad placement, reports an average lift rate of 41.23% across geotargeted campaigns⁴. Academics have confirmed the improved response rates on campaigns targeted based on the real-time geographic proximity to a retailer (e.g. Ghose, Goldfarb, and Han (2013); Luo, Andrews, Fang, and Phang (2014); Danaher, Smith, Ranasinghe, and Danaher (2015)). Danaher, Smith, Ranasinghe, and Danaher (2015) explain the appeal of mobile coupons: "They are inexpensive, quick to disseminate, and adaptable; moreover, they can convey a reasonable amount of information; appeal to notoriously difficult-to-reach younger consumers; and be customized on the basis of location, personal information, and prior purchase behavior." (page 711) They report more than 10 billion mobile coupons redeemed worldwide in 2013.

Practitioners continue to seek increasingly granular forms of *geo-precise* targeting both within a location (e.g. within a store) and across locations (e.g. distinction between "types" of locations). One increasingly popular form of geo-precise targeting is *geo-conquesting*, whereby the mobile advertiser targets customers near a competitor's location. Early practitioner reports suggest that geo-conquesting leads to even higher response rates⁵. In one of its quarterly Mobile-Location Insights Reports, Xad noted that one third of their geo-targeted campaigns now include such geo-conquesting. A recent academic study of mobile promotions for a movie theater finds that real-time targeting of mobile customers near a competitor's location can increase purchase rates, with higher incremental purchases for very deep discounts off the regular price (Fong, Fang, and Luo (2015)).

In addition to location, mobile advertisers can also target customers based on behavior. The combination of geographic and behavioral targeting should enable firms to triangulate on the most geographically relevant customers. In our study, we combine geographic location with historic visit behavior. We use the recency measure to proxy for differences in willingness-to-pay across customers within a location.

A potential limitation of the existing body of evidence for mobile targeting is the omission of strategic

⁴"Rocket Fuel Proves Digital Ads Influence Physical Activity, Drives 41.3% Lift in Store Visits With Programmatic Targeting," *Business Wire*, February 17, 2015.

⁵Mark Walsh, "Geo-Conquesting' Drives Higher Mobile Click Rates," *Online Media Daily*, May 17, 2013. [Accessed on 11-7-2015 at <http://www.mediapost.com/publications/article/200578/geo-conquesting-drives-higher-mobile-click-rates.html?edition=>].

considerations. The evidence typically studies the incentives for a single firm to geo-target offers, holding competitor actions fixed. For instance, YP Marketing Solutions recently ran a hyper geo-targeted mobile campaign for Dunkin' Donuts that "targeted competitors' customers with tailored mobile coupons."⁶ They reported a 3.6% redemption rate among mobile users that clicked and took secondary actions. However, this analysis held competitors' actions fixed. In other words, the existing evidence studies targeted marketing through the lense of a monopoly theory. Given the evidence of a strong incentive for a focal firm to poach its competitor's customers, one might expect the competitor to face symmetric incentives to implement geo-conquesting campaigns. Our work contributes to the literature by analyzing the returns to geo-conquesting in a competitive environment. We design a large-scale field experiment that allows us to analyze geo-conquesting through the lense of oligopoly theory rather than through the lense of monopoly theory. Our findings indicate that the returns to geo-conquesting may be overstated when equilibrium considerations are ignored.

3 Field Experiment

3.1 Experimental Design

Our data come from a unique field experiment conducted with the cooperation of a large, wireless service provider that provides the platform for targeted mobile promotions. In the experiment, a mobile SMS promotion consisted of an offer to buy one general admission voucher for any 2D movie showing at a given movie theater on the day the SMS was sent. The SMS contained a brief description of the offer, and recipients could click on a link to purchase the voucher and take advantage of the price discount for any movie showing in the theater on that day. In practice, an advertiser pays 0.08 RMB per message sent. Since this study was coordinated with the wireless provider, all of the messages used in our campaigns were paid by the wireless provider, not the theaters⁷.

Our subject pool consists of mobile subscribers that were randomly sampled over the course of an hour on a Saturday morning between 11:00 a.m. and 12:00 noon. We used this early Saturday afternoon time slot is an off-peak time for the movie theater to ensure that we do not exceed theater capacities. The subjects were sampled from two locations: the 500-meter-radius geofences surrounding two competing theaters (hereafter referred to as firms A and B), located in two large shopping centers respectively. Each subject was classified into one of four segments based on her observed geographic location and type. The geographic location represents the mall in which the subject was located, A or B, at the time of

⁶"Mobile Retargeting, Optimization & Hitting the ROI Bullseye," accessed on YP Marketing Solutions website on 12/1/2015 at http://national.yip.com/downloads/MMS_Chicago.pdf.

⁷In this experiment, we do not vary the price per message sent, treating it as exogenous. We can derive each theater's incremental revenues per message sent and, hence, demand for messaging services. However, the set of customers for SMS services spans a much broader range of markets than theaters, such as gaming, apps, call services, restaurants, travel, education, and mobile news. Each of these markets likely has different SMS service demand due to differences in the degree of competition and the magnitude of incremental revenue potential. In sum, our data are not suitable for studying the mobile platform's pricing incentives for access to SMS messaging.

the intervention. Based on the location, a consumer's type was then determined by her historical visits to the theater in the corresponding mall, a measure of recency. A movie theater visit was defined by any dwell time in the theater of at least 90 consecutive minutes, measured using the GPS and cell tower triangulation on mobile signal reception for the individual's phone. A consumer was classified as "high type" if she visited the movie theater at least once during the preceding two months; otherwise she was classified as "low type." None of our subjects previously received SMS promotions from either theater. Therefore, the "high" versus "low" classification is based on visiting the theater at regular box-office prices. The four consumer segments therefore consist of: (A, High), (A, Low), (B, High), and (B, Low).

Each subject was randomly assigned to one of several promotional conditions based on a 3x3 design. Based on a subject's location, the local theater was classified as "defensive" and the more distant theater was classified as "offensive." We use a symmetric design such that we can analyze each of theaters A and B from the offensive and defensive perspective. In the control condition, a subject did not receive an SMS offer. In our SMS promotion conditions, the discount depths were chosen based on the mobile carrier's experience with previous promotions and based on the pilot study in Fong, Fang, and Luo (2015). In our "defense only" condition, a subject received an SMS from the defensive theater reading: "To buy a voucher for general admission to any of today's 2D showings at [defensive theater] at a [20%, 40%] discount, follow this link: _." In our "offense only" condition, the subject received an SMS from the offensive theater reading: "To buy a voucher for general admission to any of today's 2D showings at [offensive theater] at a [40%, 60%] discount, follow this link: _." In our "offense and defense" condition, the subject received two SMSs, one from each of the offensive and defensive theaters respectively. In the promotion cells, we directly observed whether or not a subject purchased a voucher through the SMS offer.

To construct the sample, we began by sampling mobile subscribers located in the two shopping malls' respective geofences at the time of the study. During the course of the hour when subjects were sampled, approximately 57,000 mobile subscribers were observed across the two locations. The population weights associated with each of the four observed consumer segments are: 12% (A, High), 26% (A, Low), 18% (B, High) and 44% (B, Low). With 9 pricing conditions applied to each of the four observed segments, the experimental design involves a total of 36 cells, or 9 cells per segment. Approximately 500 subjects were assigned to each cell, with a total sample size of 18,000 subscribers. Within each consumer segment, the randomization of subjects across the 9 pricing cells was performed by assigning each sampled mobile subscriber a random uniform integer between 1 and 9. We counter-balanced the order in which the offers from each of the two theaters were received for those subjects receiving an SMS from each theater. Testing for sequential promotions was outside the scope of this study.

For each subject, we observe whether or not a ticket was purchased from one of the movie theaters. We directly observe when a subject purchased one of the movie vouchers offered in a promotional SMS. To determine the rate at which our control subjects bought movie tickets in the "no promotion" cells, we use the GPS and cell tower triangulation on mobile signal reception for a subject's mobile device to track whether the subject visited one of our two theaters and dwelled in the theater's premises for

at least 90 consecutive minutes. To control for the possible purchase acceleration associated with our time-sensitive SMS offers, we tracked the control subjects for eleven days to assess whether they “ever” went to a movie at one of the two theaters. We used an eleven-day period to avoid overlapping with the timing of release of new movies that could change demand.

In total, 553 of our 18,000 subscribers purchased a ticket representing about 3.1% of the sample. 535 of the 16,000 subjects receiving at least one SMS offer purchased a voucher, a promotional response rate of 3.3%. This 3.3% conversion rate is comparable to other recent targeted pricing experiments on mobile phones (e.g. Dubé, Luo, and Fang (2015)). We never observe a subject visiting a movie theater more than once in the “no promotion” control case, nor do we observe a subject purchasing more than one movie voucher in response to the SMS offer. Therefore, we can treat consumer demand as discrete choice.

In addition to observing the exact prices charged and purchase decisions of each subject, we also observe several measures of each subject’s mobile usage behavior. In particular, we observe each subject’s average monthly phone bill (ARPU), total minutes used (MOU), short message services (SMS), and data usage (GPRS). Summary statistics for the mobile usage variables are reported in Appendix B, Table 9 by segment.

We report randomization checks in Appendix B, Table 10. Making all pairwise comparisons for the 9 cells (36 comparisons) for our 4 mobile usage variables (for a total of 144 comparisons) resulted in 6 differences in means at the .05 significance level, and an additional 4 differences at the .10 significance level, which is not a rate greater than chance. Running a Tukey Test for each mobile usage variable to correct for multiple comparisons, no significant differences were found between any pair of cells, for any of the 4 mobile usage variables. In Appendix B, Table 11, we also provide summary statistics about each of the two malls, finding that both cater to similar demographic profiles of customers.

4 Experimental Results

An advantage of the design of the experiment is that we can analyze some aspects of the promotional effects model-free. Basic tests for differences in aggregate purchasing across pricing conditions are provided in Appendix B, Tables 12 and 13. The analogous tests are reported by consumer segment in Table 14 to 17.

For our model-free analysis of uniform pricing across segments, we collapse the data across consumer type segments and pool the two firms into “defensive” and “offensive” states. The corresponding purchase rates by promotion condition are reported in Figure 1. Moving down the first column, we see that defensive promotions raise demand substantially. Increasing the discount from 20% to 40% doubles sales from 2.3% to 5.2% (p-value<.01)⁸. Therefore, as expected, a theater faces a downward-sloping

⁸All p-values reported in this section are derived from OLS regressions of purchase indicator on the full set of experimental conditions, and obtain the same results using conventional or robust standard errors (we report whichever is larger). Alternative methods of testing differences in mean purchase rates are described in the Appendix, Table 12.

demand curve in its local market. To a lesser extent, offensive promotions also increase sales, as seen by moving along the first row. None of the sample consumers purchase a ticket from the offensive firm at regular price. While a 40% discount does generate offensive ticket sales, the level is quite low. However, increasing the offensive discount from 40% to 60% increases offensive ticket sales substantially from 0.35% to 2.8% (p-value<.01)⁹. Interestingly, these results confirm that a theater faces a downward-sloping demand curve in its competitor's local market and that the "other mall" represents a potential market for a theater. Moreover, the differences in defensive and offensive price sensitivities suggest an opportunity for geographic price targeting.

The remaining cells of the experiment measure the cross-promotional effects. Interestingly, defensive promotions appear to be immune to an offensive promotion. As we move along the columns of the second and third row, we fail to reject that the level of sales for the defensive firm remains unchanged and the cross-effects of offensive promotions on defensive demand are all statistically insignificant. If we focus on the high segment in location A only (see Figure 3), then we do see a slight cross-effect from the offensive promotion. The demand for defensive tickets at a 40% discount falls from 5.95% to 4.92% when the offensive promotion is increased from 40% to 60%, though this difference is not significant (p-value=.47). Therefore, demand for the "offensive" theater comes primarily from the outside good – i.e. the conversion of non-purchasers conditional on the defensive firm's prices. This role for category expansion indicates that the full market coverage assumption in Shaffer and Zhang (1995) does not apply in our study of movie theaters and that a prisoner's dilemma is therefore not a foregone conclusion.

The effects of the offensive promotions do appear to be mitigated by a defensive promotion. As we move down the rows of the third column, we observe a large and statistically significant decline in the level of offensive ticket sales compared to the response when there is no defensive promotion. For instance, increasing the offensive discount from 40% to 60% increases offensive ticket sales by 2.5% (p-value<.01) when there is no defensive discount, but only by 0.70% when there is a defensive discount of 20% (p-value=.010, difference-in-differences p-value<.01). Therefore, the defensive theater not only draws its demand from converting local non-purchasers, it also draws demand away from the offensive theater. Substitution patterns between the two firms are therefore asymmetric, with the offensive firm facing more competition from the defensive firm than vice versa. Our analysis of cross-promotion effects indicates that both firms compete in their respective offensive and defensive markets. Therefore, the intuition from monopoly price discrimination theory is unlikely to provide a reliable analogy for price-targeting in our study.

We now turn to each firm's pricing incentives. The average revenues per potential consumer (as a proportion of the full ticket price, net of discounts) are reported in Figure 2. The incremental purchases from the promotions generate incremental revenues per customer for both the offensive and defensive firm. Furthermore, the cross-promotional effects on revenues are asymmetric. Both firms, however, have a clear incentive to implement the deep discount. Since the equilibrium uniform prices could lie either

⁹This finding is consistent with Fong, Fang, and Luo (2015).

above or below the deep discount levels, it is possible that we do not observe each firm’s best-responses under either uniform or geographically targeted pricing.

We now break the purchases and revenues apart for each of the four consumer segments; although we still pool the two firms into “defensive” and “offensive” states as above. The corresponding purchase rates by promotion condition are reported for each segment in Figure 3. Most of our intuition about the sales lift from discounts in the uniform pricing case above carry over to the segmented pricing case. There is an asymmetric cross-promotional effect in each segment whereby the offensive firm’s promotions are offset by the defensive firm, but not vice versa. The sales levels are smaller in the Low segment than in the High segment, as expected. Therefore, both firms compete in the two consumer type markets, although competition appears to be more intense in the “High” market than in the “Low” market.

An interesting feature of the experimental data is that none of the low consumer types purchases a ticket at the regular price. However, aggressive discounts are capable of drawing a substantial number of these consumers to buy tickets. In both locations, a defensive discount of 40% attracts over 4% of the low consumer segment to purchase.

Looking once again at pricing incentives, we report the average revenues per potential consumer by segment in Figure 4. As in the uniform case, both firms appear to have a strong incentive to offer deep discounts in each of the consumer segments. We therefore have strong evidence that firms will want to implement discounts. However, neither the optimal unilaterally targeted prices nor the equilibrium targeted prices are likely observed within the discrete set of tested price levels. We overcome this limitation of the experiment by using a demand model to predict outcomes for a continuum of price combinations.

In sum, our model-free analysis indicates that each firm appears to have an incentive to use SMS discounts in each of the consumer segments. However, the returns to a targeted SMS campaign are clearly mitigated by the targeting efforts of a competitor. We also observe a stark geographic asymmetry whereby the offensive firm’s targeting efforts are more vulnerable to those of a defensive firm than vice versa. However, the coarseness of the price grid used in the experiment prevents us from observing each firm’s best-response and, hence, from observing the likely equilibrium prices that would likely prevail in this market.

5 Probit Demand and Estimation

We need a model to resolve the fact that the equilibrium prices lie below the range of prices tested in our experiment. In this section, we describe the trinomial probit model of demand that we use to estimate demand along the entire support of prices. By allowing for correlated errors, the model should be flexible enough to fit the observed patterns of asymmetric cross-promotional effects for the defensive versus the offensive firm in a market described in section 4. Another popular specification is the random coefficients logit, which is easier to estimate. The flexibility of this model would come from measuring within-subject, unobserved heterogeneity in tastes for theaters. Since we have cross-sectional data in

this application, we prefer the probit which allows for flexible substitution patterns without the need for explicitly modeling within-subject unobserved heterogeneity.

5.1 Consumer Demand

Let $h = 1, \dots, H$ denote the consumers and $j = A, B$ denote the theater alternatives, where $j = C$ is the no-purchase alternative. Each consumer belongs to one of the $k = 1, \dots, K$ observable segments (based on geography and historic purchase intensity). At the population level, the segment proportions are denoted by $\{\lambda^k\}_{k=1}^K$. Each theater has attributes $X_j = (\mathbb{I}_{1j}, \mathbb{I}_{2j}, p_j)$ where $\mathbb{I}_{1j} = 1$ if j is theater 1 and 0 otherwise. p_j is the ticket price at theater j .

Assume a consumer h in segment k obtains the following indirect utility from her choice:

$$u_{hA} = \theta_A^k - \alpha^k p_A + \tilde{\epsilon}_{hA}$$

$$u_{hB} = \theta_B^k - \alpha^k p_B + \tilde{\epsilon}_{hB}$$

$$u_{hC} = \tilde{\epsilon}_{hC}$$

where $\tilde{\epsilon}$ are random utility shocks. Let $\eta_h \equiv \begin{bmatrix} \tilde{\epsilon}_{hA} - \tilde{\epsilon}_{hC} \\ \tilde{\epsilon}_{hB} - \tilde{\epsilon}_{hC} \end{bmatrix} \sim N(0, \Psi^k)$ if consumer h is in segment k . We can write the vector of theater-specific indirect utilities as $U_h \equiv \begin{bmatrix} u_{hA} \\ u_{hB} \end{bmatrix} = B^k X + \eta_h$, where $X = \begin{bmatrix} X'_A \\ X'_B \end{bmatrix}$ is (6×1) vector of product attributes and $B^k = I_2 \otimes \beta^{kT}$ is a (2×6) matrix of parameters with $\beta^k = (\theta_A^k, \theta_B^k, \alpha^k)$. Finally, index consumer choices by $y_h \in \{A, B, C\}$.

The expected probability that a consumer h in segment k chooses alternative j is

$$Pr(y_h = j | B^k, X, \Psi) = Pr(u_{hj} - u_{hi} > 0, \forall i \neq j).$$

We can simplify our analysis by using the following transformations for consumer h in segment k :

$$Z^{k,(A)} = \begin{bmatrix} u_{hA} - u_{hB} \\ u_{hA} \end{bmatrix}, Z^{k,(B)} = \begin{bmatrix} u_{hB} - u_{hA} \\ u_{hB} \end{bmatrix}$$

or in matrix form

$$Z^{k,(j)} = \Delta^{(j)} U_h, j \in \{A, B\}$$

where

$$\Delta^{(A)} = \begin{bmatrix} 1 & -1 \\ 1 & 0 \end{bmatrix}, \Delta^{(B)} = \begin{bmatrix} -1 & 1 \\ 0 & 1 \end{bmatrix}$$

and $E\left(Z^{k,(j)}\right) \equiv \mu_Z^{k,(j)} = \Delta X B^k$, $Var\left(Z^{k,(j)}\right) \equiv \Sigma_Z^{(j)} = \Delta^{(j)} \Psi^k \Delta^{(j)T}$ and $corr\left(Z^{k,(j)}\right) \equiv \rho_Z^{(j)} = \frac{\Sigma_{Z(A,B)}^{(j)}}{\sqrt{\Sigma_{Z(A,A)}^{(j)} \Sigma_{Z(B,B)}^{(j)}}}$.

If we standardize $Z^{k,(j)}$, we obtain $z^{k,(j)} = \left[D^{(j)}\right]^{-\frac{1}{2}} Z_h^{k,(j)}$, where $D^{(j)} = \text{diag}\left(\Sigma_Z^{(j)}\right)$ and $E\left(z^{k,(j)}\right) \equiv \mu_z^{k,(j)} = \left[D^{(j)}\right]^{-\frac{1}{2}} \Delta X B^k$. We can now write the choice probabilities as follows:

$$\begin{aligned} Pr\left(y = j | B^k, X, \Psi^k\right) &= Pr\left(z^{k,(j)} > 0 | \mu_z^{k,(j)}, \rho_Z^{(j)}\right) \\ &= \int_{-\infty}^{\mu_{zA}^{(j)}} \int_{-\infty}^{\mu_{zB}^{(j)}} \phi\left(x, y; \rho_Z^{(j)}\right) dy dx, \quad j \in \{A, B\} \\ &= \Phi\left(\mu_{zA}^{(j)}, \mu_{zB}^{(j)}; \rho_Z^{(j)}\right) \end{aligned} \quad (1)$$

where $\phi\left(x, y; \rho\right) = \frac{1}{2\pi\sqrt{1-\rho^2}} \exp\left(-\frac{x^2-2xy\rho+y^2}{2(1-\rho^2)}\right)$ and $\Phi\left(x, y; \rho\right) = \int_{-\infty}^x \int_{-\infty}^y \phi\left(u, v; \rho\right) dv du$, and $Pr\left(y = C | B^k, X, \Psi^k\right) = 1 - \sum_{j \in \{A, B\}} Pr\left(y = j | B^k, X, \Psi^k\right)$.

The probabilities in equation (1) give rise to the usual trinomial probit model of choice. We estimate these choice probabilities using the MCMC algorithm proposed by McCulloch and Rossi (1994). The following priors are assigned

$$\begin{aligned} B^k &\sim N\left(\bar{B}, A^{-1}\right) \\ \Psi^k &\sim IW\left(\nu, V\right) \end{aligned}$$

where $A = \text{diag}(0.01)$, $\nu = 6$ and $V = \text{diag}(\nu)$. The estimation algorithm is defined over the unidentified parameter space. In our results, we report posterior distributions for the identified quantities. We conduct all of our analysis in R, using the `rmnpGibbs` function from the `bayesm` package. Since we expect the taste coefficients and random utility covariances to differ across segments, we estimate the model separately for each segment. For each run of the estimation algorithm, we simulate a chain with 200,000 posterior draws and retain the last 100,000 draws for inference.

5.2 Aggregate Demand and Substitution Patterns

To derive the demand system for each theater, we integrate over all the consumers in the population. The total market share for theater j is

$$S_j(p) = \sum_k \lambda^k Pr\left(y = j | B^k, p, \Psi^k\right) \quad (2)$$

where we focus on the price vector $p = (p_A, p_B)$ and drop the theater dummy variables to simplify the notation hereafter.

The own- and cross-price elasticities of the total market share for theater j are:

$$\varepsilon_{jj} = \frac{p_j}{S_j} \sum_k \lambda^k \frac{\partial Pr(y = j | B^k, p, \Psi)}{\partial p_j}$$

and

$$\varepsilon_{ji} = \frac{p_i}{S_j} \sum_k \lambda^k \frac{\partial Pr(y = j | B^k, X, \Psi)}{\partial p_i}$$

respectively. Exact expressions for the derivatives, $\frac{\partial Pr(y=j|B^k,X,\Psi)}{\partial p_i}$, are derived in Appendix A.

6 Pricing

We now analyze several different Bertrand-Nash pricing scenarios between the two competing theaters. A novel feature of our analysis is that we study the oligopoly pricing implications in a market with probit demand. An advantage of the probit model is that it does not exhibit the IIA property, a characteristic of logit demand systems which can potentially lead to unrealistic substitution patterns and pricing patterns. For instance, under logit oligopoly, firms' equilibrium prices, mark-ups and pass-through rates are proportional to their market shares. Like many oligopoly models with empirically realistic demand specifications (e.g. random coefficients logit), it is not possible to prove existence and uniqueness of a Bertrand-Nash price equilibrium for a probit demand oligopoly except under very strong independence assumptions (Mizuno (2003))¹⁰. In this section, we derive the the first-order necessary conditions for each of the pricing scenarios analyzed. In our numerical simulations, existence is established by computing a fixed-point to the system of necessary conditions. Uniqueness is verified by inspecting each firm's best-response function over a reasonable range of prices (0 RMB to 75 RMB in our application).

6.1 Uniform Pricing

Under uniform pricing, each firm sets a single price across all consumer segments. We study the data-based decision problem for each firm, using the posterior expected profits as the firm's reward function. Posterior profits are computed based on the posterior distribution of demand, which we simulate using the R posterior draws from the chain used to estimate the demand function via our MCMC algorithm. Implicitly, we assume firms are risk neutral and form the following posterior beliefs about demand conditional on the data, \mathbf{D} : $\left(\{B^{r,1}, \Psi^{r,1}\}_{r=1}^R, \dots, \{B^{r,K}, \Psi^{r,K}\}_{r=1}^R \right)$.

¹⁰In addition to the correlated errors, our demand specification differs from the standard independent probit because it is a discrete mixture of the probit demands in the two geographic and two consumer type segments.

Firm j 's optimal uniform pricing problem consists of picking the price $p_j^{uniform}$ as follows

$$\begin{aligned} p_j^{uniform} &= \underset{p_j}{\operatorname{argmax}} \{ p_j \sum_{k=1}^K \lambda^k \mathbb{E} [Pr(j|B^k, p, \Psi^k) | \mathbf{D}^k] \} \\ &\approx \underset{p_j}{\operatorname{argmax}} \{ p_j [\sum_{k=1}^K \lambda^k \frac{1}{R} \sum_{r=1}^R Pr(j|B^{r,k}, p, \Psi^{r,k})] \} \end{aligned} \quad (3)$$

which generates the following first order necessary conditions

$$\sum_{k=1}^K \lambda^k \sum_{r=1}^R Pr(j|B^{r,k}, p, \Psi^{r,k}) + p_j^{uniform} \sum_{k=1}^K \sum_{r=1}^R \lambda^k \frac{\partial Pr(j|B^{r,k}, p, \Psi^{r,k})}{\partial p_j} = 0. \quad (4)$$

Please see Appendix A for the derivation of the slopes of the probit demand system. Firm j can assess the choice of $p_j^{uniform}$ by studying the corresponding posterior distribution of profits:

$$\left\{ p_j^{uniform} \sum_{k=1}^K \lambda^k \frac{1}{R} Pr(j|B^{r,k}, p, \Psi^{r,k}) \right\}_{r=1}^R.$$

A Nash equilibrium in uniform prices consists of a vector of uniform prices for each firm, $p = (p_A, p_B)'$ for which each prices satisfies its corresponding set of first order necessary conditions as in equation 4.

6.2 Targeted Pricing

The mobile platform can potentially enable different types of targeting across our $K = 4$ segments: location, type, location & type. Suppose firms are able to target the groups in Ω , a partition of the $K = 4$ segments. For instance, geographic targeting would consist of

$$\begin{aligned} \Omega &= \{ \text{location A, location B} \} \\ &\quad \{ \{ \text{high type in location A, low type in location A} \}, \{ \text{high type in location B, low type in location B} \} \}; \end{aligned}$$

type targeting would consist of

$$\begin{aligned} \Omega &= \{ \text{high type, low type} \} \\ &\quad \{ \{ \text{high type in location A, High type in location B} \}, \{ \text{low type in location A, low type in location B} \} \}; \end{aligned}$$

and the mixed targeting would consist of

$$\Omega = \{ \text{high type in location A, low type in location A, high type in location B, low type in location B} \}.$$

Given targetability across the elements of a partition, $\omega \in \Omega$, firm j 's targeted pricing problem consists of picking the vector of prices p_j^Ω as follows

$$\begin{aligned}
p_j^\Omega &= \underset{p_j}{\operatorname{argmax}} \left\{ \sum_{\omega \in \Omega} p_{j\omega} \sum_{k \in \omega} \lambda^k \mathbb{E} [Pr(j|B^k, p_\omega, \Psi^{r,k}) | \mathbf{D}^k] \right\} \\
&\approx \underset{p_j}{\operatorname{argmax}} \left\{ \sum_{\omega \in \Omega} p_{j\omega} \sum_{k \in \omega} \lambda^k \frac{1}{R} \sum_{r=1}^R Pr(j|B^{r,k}, p_\omega, \Psi^{r,k}) \right\}
\end{aligned} \tag{5}$$

which generates the following first order necessary conditions

$$\sum_{k \in \omega} \left(\lambda^k \sum_{r=1}^R Pr(j|B^{r,k}, p_\omega, \Psi^{r,k}) + p_{j\omega}^\Omega \sum_{r=1}^R \lambda^k \frac{\partial Pr(j|B^{r,k}, p_\omega, \Psi^{r,k})}{\partial p_j} \right) = 0, \forall \omega \in \Omega \tag{6}$$

Please see Appendix A for the derivation of the slopes of the probit demand system. Firm j can assess the choice of p_j^Ω by studying the corresponding posterior distribution of profits:

$$\left\{ \sum_{\omega \in \Omega} p_{j\omega}^\Omega \sum_{k \in \omega} \lambda^k \frac{1}{R} Pr(j|B^{r,k}, p, \Psi^{r,k}) \right\}_{r=1}^R$$

A Nash equilibrium in targeted prices over consumer partition C consists of a vector of targeted prices for each firm, $p^\Omega = (p_A^\Omega, p_B^\Omega)$ for which each price satisfies its corresponding set of first order necessary conditions as in equation 6.

A firm that unilaterally tests targeted pricing effectively assumes its rival would not deviate from current pricing. Suppose firm 1 unilaterally targets prices across the segment partition C . We define the corresponding unilateral targeting as the scenario in which firm 1 solves the monopoly price discrimination problem, leading to a price vector $p^{\Omega, \text{unilat}} = (p_1^\Omega, p_2^{\text{uniform}})$ with p_1^Ω and p_2^{uniform} defined as in equations 3 and 5 respectively.

7 Results

7.1 Demand Estimates

We now discuss our demand estimates. A unique feature of the demand estimation exercise is that the prices consumers face at each theater were generated through a randomization, eliminating the usual endogeneity concerns associated with observational marketing data¹¹. Recall that the pool of consumer subjects come from four distinct segments based on their geographic location and historic movie theater visit behavior: (i) high consumers in location A, (ii) low consumers in location A, (iii) high consumers in location B and (iv) low consumers in location B. We estimate a separate choice model in each of the four segments. As we explain below, we want the theater-specific intercepts and the covariance terms to be segment-specific. The intercepts will help us fit differences in response rate levels across segments. The covariance terms will help us fit the non-IIA substitution behavior differences across cells. We use both

¹¹For instance, Chintagunta, Dubé, and Goh (2005) simulate monopoly targeting using household-level estimates obtained from an instrumental variables estimator to resolve the endogeneity in observational price variation.

a multinomial logit model and a multinomial probit model to verify whether the IIA problem from the former leads to inferior fit. Recall that we do not need to estimate the segment weights, λ^k , as these are observed by matching mobile subscribers with their location at the time of the study and their historic theater visit behavior.

Table 1 summarizes the posterior fit of each specification, by segment, using the posterior log marginal likelihood computed with the method of Newton and Raftery (1994). In each case, we trimmed the lower and upper one percent of draws to avoid underflow problems¹². The results indicate that the additional flexibility of the multinomial probit, which allows for correlated and heteroskedastic random utility shocks, improves fit in the high consumer segments in both locations. However, the probit only fares marginally better than the logit in the low consumer segments in both locations. The improved fit of the probit stems largely from its ability to fit the cross-promotional effects, especially in the high consumer segments. In Appendix B, Figure 11 reports the offensive purchase rates in each of our experimental cells. Focusing on the high consumer segment, the probit fits the purchase rates better as the offensive firm increases its discount from 40% to 60% off the regular price. The probit allows for more flexible substitution between the two theaters relative to the outside good. As a result, offensive promotions can increase offensive ticket sales without stealing “too much” business from the defensive firm.

To assess the differences in fit of the logit and probit specifications, we report the predicted purchase probabilities along with the observed purchase rates in Figures 11 and 12. In the tables we can see that the probit does a better job predicting the cross-promotional effects noted in Section 4. Recall that for both theater locations, the offensive purchase rate in the high consumer segments is sensitive to the defensive price. The probit predictions better capture this sensitivity. Thus, we conclude that the added flexibility of the probit is better suited for modeling demand in this market.

Hereafter, we focus our results on the multinomial probit specification. We report the estimated coefficients from the multinomial probit in Table 2. For each coefficient, we report the posterior mean and the 90% posterior credibility interval. We observe quite a bit of heterogeneity across the consumer segments, as one would expect. Most important, we find a lot of heterogeneity in the distribution of the utility shocks, especially with regards to the covariance in the shocks. Our point estimates are consistent with substantial heteroskedasticity, although we cannot rule out that the variance of theater B shocks is one. At the bottom of the table, we report the correlations, $\rho_{A,B}$. We find that the theater-specific utility shocks are highly positively correlated for the high consumer types in both locations. The strong evidence for correlation explains why we select the multinomial probit in favor of the multinomial logit in each segment. The intuition can be seen in the raw data. All offensive discounts in the high segment draw some demand away from the defensive firm. However, only the large discount of 60% off draws new buyers into the category. For the low segment, the theater-specific shocks are highly negatively correlated in location A. In our data, we observe almost no substitution between the theaters in this segment, consistent

¹²The trimming avoids very small-valued draws that could lead to numerical problems with the calculation of the harmonic mean.

with strong idiosyncratic preference for a specific theater. Accordingly, any increases in demand from a discount appear to derive from category expansion. The correlation is small and positive in segment B and, once we account for parameter uncertainty, we cannot rule out that the correlation is zero. The relatively high correlations in the strong segments will also intensify demand in these segments.

Next, we turn to our estimated price elasticities. In Table 3 we report the posterior mean own and cross-price elasticities in each consumer segment, computed at the regular prices (both charge 75 RMB) and also at the largest discount of 60% (both charge 30RMB). As expected, the low consumer segment has higher own elasticities than the high consumer segment at both price levels. In all four segments, we find that both firms' regular prices of 75RMB are at very elastic regions of their respective demand curves. Given that both theaters are far from capacity at regular prices during the off-peak time slot¹³, we would expect the effective marginal cost per ticket to be close to zero. In our analysis, we ignore the potential role of concession revenues. Hence, optimized uniform pricing should be at the unit-elastic region of total demand, which is the weighted average of the segment demands. If firms are optimizing their profits and setting uniform prices, they should be operating on the inelastic region of at least one of the segments. This evidence suggests an opportunity for firms to generate substantial demand during off-peak hours through large discounts off their box-office prices, which are uniform across all time slots (peak and off-peak). This result is consistent with the substantial returns to large discounts we observed in each segment in our raw experimental data as discussed in section 4 above and in figures 2 and 4.

We can see these patterns by looking at the estimated demand functions plotted in Figures 5 and 6. Each plot reports the posterior expected demand function along with the 90% posterior credibility interval. In Figure 5 we can see the asymmetric substitution patterns. When Theater B lowers its price, the shift in demand for Theater A in Mall A is minimal. However, the shift in demand for Theater A in Mall B is quite large. Figure 5 illustrates the much higher intensity of competition in the High segment than in the Low segment. In Mall B, a decrease in the price for Theater B has a much larger effect on High demand for Theater A than Low demand for Theater A.

The cross elasticities are also as expected. In the high consumer segments, we observe highly asymmetric cross-elasticities, with the offensive firm's demand much more vulnerable to the defensive firm's price. In the low consumer segments, we observe relatively little substitution between the two theaters, meaning that discounts mostly draw new consumers into the category.

7.2 Uniform Pricing

Since firms do not appear to be optimizing their uniform (across all time slots) box-office prices to the mid-day market in our analysis, it is difficult to discern in the raw data whether the large returns to discounting are due to optimization for the time slot or to targeting. Since the optimal uniform prices

¹³Theater A has a capacity of 1,200 seats per show and theater B has a capacity of 2,000 seats per show. With 4 shows per day and mobile subscribers representing 75% of total mall traffic on a typical Saturday afternoon, the theaters are less than half full for the average show.

cannot be inferred from the raw data, we instead use our trinomial probit demand estimates to compute them. We combine our demand estimates, from section 7.1, with the system of first order necessary conditions, equation (4). Results are displayed in Table 4. Consistent with the experimental results, both firms have strong incentives to reduce their prices. In equilibrium, Firm A charges 19.29 RMB per ticket and firm B charges 18.86 RMB per ticket, drawing in substantial demand, especially from the low consumer segment.

7.3 Unilateral Targeting Results

Having established the benefits of optimizing their SMS prices, we can now explore each theater's incentives to target prices across the different consumer segments. In practice, most firms are unable to manipulate their competitors' prices experimentally. A more realistic scenario involves a firm field testing its own targeting opportunities while holding its competitor's prices fixed. In essence, the typical field test applies monopoly price discrimination theory. We study this scenario by allowing one firm to optimize its segment-specific prices under the assumption that the other firm does not deviate from its uniform Nash equilibrium price. An interpretation of this scenario is that the competitor does not detect the focal firm's deviation from Nash pricing during the field experiment.

We use our probit demand estimates to compute the optimal price for a firm that unilaterally changes its price and disregards strategic considerations. The targeting firm sets prices to satisfy the optimality conditions in equation 6 evaluated at the competitor's Nash equilibrium price computed in the previous section. In several of our targeting schemes, demand consists of a mixture over different consumer types. Under geographic targeting demand for theater j in location k is given by:

$$S_{j|k} = \lambda^{(low,k)} Pr\left(y = j | B^{(low,k)}, p_k, \Psi^{(low,k)}\right) + \lambda^{(high,k)} Pr\left(y = j | B^{(high,k)}, p_k, \Psi^{(low,k)}\right).$$

Similarly, under type targeting, demand for theater j from type k customers is given by:

$$S_{j|k} = \lambda^{(A,k)} Pr\left(y = j | B^{(A,k)}, p_k, \Psi^{(A,k)}\right) + \lambda^{(B,k)} Pr\left(y = j | B^{(B,k)}, p_k, \Psi^{(A,k)}\right).$$

Results from the optimized unilateral targeting are reported in Table 5. The first row repeats the last row of Table 4, indicating the expected revenues per messaged consumer under uniform pricing, which we use as a benchmark. The subsequent rows report each firm's expected revenues when it unilaterally deviates from the uniform Bertrand-Nash price and charges its optimal targeted price for each of the targeting scenarios. For each firm, unilaterally targeting prices unambiguously raises expected profits in each scenario as one would expect by construction under monopoly price discrimination. Figure 7 plots the posterior distribution of the percentage difference in revenues under targeting and uniform pricing for each scenario. Less than 4% of the posterior probability mass falls in the negative region of the percentage change in profits from unilateral targeting versus uniform pricing for each scenario. Theater A generates an expected gain of 1.18% with geographic targeting, but only 0.45% under type

targeting. Targeting on type and location generates an expected gain of 1.83%. Theater B gains more from unilateral targeting with an expected gain of 3.82% under geographic targeting, 0.94% under type targeting and 4.51% under both.

We included a passive competitor in this analysis to mimic what is implicitly assumed when a firm applies monopoly price discrimination theory to study its targeting incentives. While we do not report the passive competitor's results in the Tables and exhibits, the findings are as one might expect. Under geographic targeting, the passive competitor's expected profits fall by 1% for both theaters. These losses reflect the fact that the targeting firm charges substantially lower prices in the passive competitor's local market. Interestingly, the passive firm always benefits slightly from a competitor unilaterally type targeting. This is because, under type targeting, the the targeting firm raises its price in the high consumer segment, causing some customers to substitute to the passive competitor. This incremental revenue raises total passive competitor profits.

In the last row of Table 5, we look at the combination of pure and location targeting by allowing firms to price discriminate unilaterally across all four consumer segments. As expected from monopoly price discrimination theory, both firms are unambiguously better off with this finer degree of price discrimination.

From a practitioner point of view, we can now assess the returns to targeting when we ignore strategic considerations. Relative to uniform pricing, theaters A and B unilaterally generate a 1.18% and 3.82% return respectively under geographic targeting. For unilateral type targeting, theaters A and B generate a 0.45% and 0.94% return respectively. When theaters A and B unilaterally target on both geography and type, they generate a 1.83% return and 4.51% return on investment respectively.

7.4 Targeting in equilibrium

We now allow for strategic considerations in the analysis of targeting. Under each targeting scenario, both firms set their prices to satisfy the optimality conditions in equation 6. From the theoretical literature on competitive price discrimination, we already know that the returns to targeting in equilibrium are not unambiguous.

We first look at each firm's best-response function in each of the targeting scenarios, as plotted in Figures 8 and 9¹⁴. We can immediately see that under geographic targeting, we have best-response asymmetry. Each firm considers its own market as the "strong" market and its rival's market as the "weak" market along the entire support. From Corts (1998), we know that the returns to targeting on firm profits are ambiguous in this case. In contrast, under consumer type targeting we have best-response symmetry. Each firm considers the "High" market to be strong and the "Low" market to be weak along the entire support. From Holmes (1989); Corts (1998); Armstrong and Vickers (2001) we know that equilibrium profits can rise in this scenario as long as competition is sufficiently intense in the "strong" market. We already know from the Table 3 that the cross-price elasticities are much larger in the High

¹⁴Each firm's best-responses are computed numerically using R's built-in "optim" function.

market than in the Low market. In fact, cross-price elasticities are nearly zero in the Low markets suggesting almost no competition.

In Table 6, we summarize each firm's equilibrium revenues under each targeting scenario. Beginning with type targeting, both firms' expected equilibrium profits are slightly higher than in their respective unilateral targeting scenarios¹⁵. The 90% credibility interval over the difference between equilibrium and unilateral profits is strictly positive for both firms. This result is consistent with the theoretical literature under best-response symmetry. Equilibrium price levels are reported in Table 7. Theaters A and B lower their prices by only 3.6% and 5.8% respectively in the Low market where competition is relatively light. In contrast, theaters A and B increase their prices by 18.9% and 26.0% respectively in the High market, where competition is relatively intense. This result is also visualized in Figure 9 where the intersection of the best-response functions in the Low market are very close to the uniform price equilibrium, whereas the best response functions in the High market intersect at substantially higher levels for both theaters. Figure 10 plots the posterior distribution of the percentage difference in revenues under targeting and uniform pricing for each scenario. Both firms strictly benefit from type targeting relative to uniform pricing.

In contrast, under location targeting, Table 6 reveals that both firms' expected equilibrium profits are lower than in either of their respective unilateral targeting scenarios. The 90% credibility interval over the difference between equilibrium and unilateral profits includes zero for both firms. Therefore, we cannot rule out that equilibrium location targeting profits are statistically indistinguishable from or worse than uniform pricing. In this scenario, with best-response asymmetry, the theory is ambiguous and the results are ultimately an empirical matter. Theaters A and B raise their prices by only 1.5% and 6.4% respectively in their defensive markets. In contrast, they lower their prices by 44% and 45.7% respectively in their offensive markets. In other words, each firm launches a massive price attack in one-another's local markets. While this does not lead to an all-out price war, it severely limits the extent to which firms can benefit from local price discrimination in a competitive environment. This result is also visualized in Figure 8 where the intersection of the best-response functions involve defensive prices that are very close to the uniform price levels, but the offensive prices are considerably lower.¹⁶ Looking at the top panel of Figure 10, we can see that over 48% of the posterior probability mass in the distribution of profit differences for geographic targeting are negative, for theater A. Just over 9% is negative for theater B.

In the last row of Table 6, we allow each firm to target on both type and location. In contrast with the unilateral case where each firm would unambiguously benefit from finer price discrimination, the equilibrium results are mixed. Theater A is better off with pure type targeting. Theater B is better

¹⁵We solve for the equilibrium prices satisfying the system of first-order conditions in each scenario using the Newton solver in the non-linear equation solver package "nleqslv" in R.

¹⁶The fact that profits do not unambiguously decrease relative to uniform pricing is different from the prisoner's dilemma finding in Shaffer and Zhang (1995). The current model differs in two ways. First, we do not assume full coverage meaning that there is an outside option that softens the profit impact of lower prices. Second, we do not allow perfect targeting in the sense that a firm cannot target a consumer based on her random utility shock.

off with pure location targeting. Moreover, Theater A's revenues under targeting in this scenario are statistically indistinguishable from revenues under uniform pricing whereas Theater B makes strictly larger revenues under targeting in this scenario.

From a practitioner point of view, we can now assess the much lower returns to targeting once we account for competitive response. Relative to uniform pricing, theaters A and B derive a 0.14% and 2.44% return on investment respectively under geographic targeting. For type targeting, theaters A and B generate a 0.95% and 1.18% return on investment respectively. When theaters A and B target on both geography and type, they generate a 0.63% return on investment and 1.86% return on investment respectively. Therefore, unlike the unilateral case, we can no longer conclude unambiguously that targeting on type generates additional ROI above and beyond targeting on geography.

An interesting empirical question is whether firms would endogenously choose to price discriminate in equilibrium. We study each form of targeting independently. Consider a two-period game in which each firm first commits to a pricing structure (targeting versus uniform), and then in the second period each firm plays its corresponding Bertrand-Nash pricing strategy. Table 8 sets up the payoff matrix associated with the 2×2 game of uniform versus targeted pricing between the two theaters. In the off-diagonal cells where only one firm target, we compute the equilibrium profits whereby the targeting firm's prices satisfy the optimality conditions in equation 6, and the non-targeting firm's prices satisfy the optimality conditions in equation 4. In each of the three scenarios, targeting is a best-response for each firm. Therefore, for our empirical setting, we would expect to observe both firms targeting. However, recall from above that the probability of losing making less money in the geographic pricing equilibrium than the uniform pricing equilibrium is 48%, for theater A, and 9% for theater B.

The non-IIA preferences in the multinomial probit demand framework play an important role in our findings. To investigate the role of IIA, we re-run our equilibrium targeting analysis with $\rho = 0$ to eliminate the correlation in preferences. Results are reported in Table 18 in Appendix B. The most striking difference from above is that targeting on geography reduces theater A's equilibrium profits. This is because setting $\rho = 0$ reduces substitutability for consumers located near theater B, making it harder for theater A to poach customers. However, setting $\rho = 0$ increases substitutability near theater A, making it easier for theater B to poach consumers. Consequently, theater A has a harder time poaching and, at the same time, needs to intensify its local defensive pricing. Although not reported, when we use the Logit demand system which exhibits the IIA property, we actually find that the strategic decision to target on geography versus uniform pricing creates a prisoner's dilemma whereby each firm targets and generates lower equilibrium profits than under uniform pricing. Recall that the Logit demand model exhibits inferior fit, based on the posterior marginal likelihood. Therefore, explicitly eliminating the IIA property with an unrestricted, multinomial probit demand is important for our conclusions about the equilibrium implications of targeting.

8 Conclusions

This study provides empirical evidence on the effectiveness of targeted pricing in a competitive market, using a mobile field experiment. Using a novel experimental design that independently varies the actual prices of two competing firms, our approach bridges the gap between applied theory and empirical work to provide several managerially relevant insights and methods. In particular, when the structure of consumer segments creates best-response symmetry and competition is tougher in the strong market, competitive price targeting can soften overall price competition, leading to higher profits than under uniform pricing. In contrast, when consumer segments lead to best-response asymmetry, competitive price targeting can toughen price competition, leading to lower prices and lower profits than under uniform pricing.

In practice, most firms test targeting strategies while holding their competitors' actions fixed. Implicitly, firms are applying the monopoly theory of price discrimination. However, the theory literature on competitive price discrimination shows that monopoly price discrimination may provide the wrong analogy for profitability. We find that firms have a strong unilateral incentive to target pricing in our mobile setting, and are not deterred by the threat of competitive response. However, competition moderates the profitability of targeted pricing. Interestingly, competition raises the profitability of behavioral targeting where firms face symmetric pricing incentives that soften price competition. In contrast, competition lowers the profitability of geographic targeting, where firms face asymmetric pricing incentives that toughen price competition. In sum, while competitive targeting does not result in lower profits per se, we do find that firms may mis-estimate the profitability of targeted pricing by disregarding competitive response.

For our study of movie theaters, a manager would conclude that the returns to behavioral-targeting generate a larger return on investment (approximately 1% for each firm) than geographic targeting in a competitive market where both firms target their prices. A manager would have reached the opposite conclusion had he/she disregarded competition. An evaluation of a unilateral targeting scheme in which the competitor does not deviate from its regular box-office pricing overestimates the returns to geographic targeting and underestimate the returns to behavioral targeting. As a rule of thumb, the degree of symmetry or asymmetry in a competitor's pricing incentives can provide guidance on the potential direction of bias in a unilateral evaluation.

Our analysis also reveals both the academic and managerial importance of the design of the experiment. We manipulated both firms' actions simultaneously. In practice, most firms have some understanding of their own profits conditional on their competitors' current prices. However, they are unlikely to have knowledge of how their optimal policies would change under counterfactual prices by their competitors. This study demonstrates the importance of strategic considerations when a firm contemplates the adoption of new targeting technologies. We address the fact that the experiment did not contain the best-response levels of each firm by using a structural model. This combination of both an experiment and a model offers a pragmatic solution to practitioners who might not be able to test "enough" price

points to observe the optimum or equilibrium in a model-free manner. In practice, if a firm was able to test “enough” price points, the equilibrium would be “observed,” simplifying the analysis considerably by obviating the need for the demand estimation and price optimization.

Our results apply to a static, simultaneous pricing game. While portions of the best-response functions are derived directly from the field experiment, our pricing results are nevertheless based on the assumptions that both firms would simultaneously play their respective static best responses. Our results could potentially change under other pricing conduct assumptions.

As we discuss in the paper, we do not address the potential endogeneity of the consumer segment definitions that can arise in a multi-period environment. In practice, as targeting draws more consumers into a theater, it endogenously changes the composition of the “recency” segments. In our application, we define recency based on consumers’ visits to the theater at regular box office prices, not based on targeted promotional prices. However, an interesting direction for future research would be to explore how dynamics affect equilibrium targeting and whether firms would continue to profit from behavioral targeting. Moreover, it would be interesting to explore whether behavioral targeting would involve targeting lower prices to firms’ strongest local customers in such a dynamic setting as in Shin and Sudhir (2010). We also assume that consumer locations are exogenous. However, another interesting direction for future research would be to explore whether consumers change their mall visiting behavior in response to their experiences with different degrees of targeted pricing across locations, as in Chen, Li, and Sun (2015).

References

- ARMSTRONG, M., AND J. VICKERS (2001): “Competitive Price Discrimination,” *Rand Journal of Economics*, 32, 1–27.
- BERRY, S. T. (1994): “Estimating Discrete-Choice Models of Product Differentiation,” *Rand Journal of Economics*, 25(2), 242–262.
- BESANKO, D., J.-P. DUBÉ, AND S. GUPTA (2003): “Competitive Price Discrimination Strategies in a Vertical Channel with Aggregate Data,” *Management Science*, 49, 1121–1138.
- BORENSTEIN, S. (1985): “Price Discrimination in Free-Entry Markets,” *Rand Journal of Economics*, 16, 380–397.
- BORENSTEIN, S., AND N. ROSE (1994): “Competition and Price Dispersion in the U.S. Airline Industry,” *Journal of Political Economy*, 102, 653–683.
- BORENSTEIN, S., AND A. SHEPARD (1994): “Competition and Price Dispersion in the U.S. Airline Industry,” *Journal of Political Economy*, 102, 653–683.
- BUSSE, M., AND M. RYSMAN (2005): “Competition and price discrimination in Yellow Pages advertising,” *Rand Journal of Economics*, 36, 378–390.
- CHEN, Y., X. LI, AND M. SUN (2015): “Competitive Mobile Targeting,” *Working Paper*.
- CHINTAGUNTA, P., J.-P. DUBÉ, AND K.-Y. GOH (2005): “Beyond the Endogeneity Bias: The Effect of Unmeasured Brand Characteristics on Household-Level Brand Choice Models,” *Management Science*, 51, 832–849.
- CORTS, K. S. (1998): “Third-degree price discrimination in oligopoly: all-out competition and strategic commitment,” *The RAND Journal of Economics*, 29, 306–323.
- DANAHER, P. J., M. S. SMITH, K. RANASINGHE, AND T. S. DANAHER (2015): “Where, When and How Long: Factors that Influence the Redemption of Mobile Phone Coupons,” *Journal of Marketing Research*, 52, 710–725.
- DUBÉ, J.-P., X. LUO, AND Z. FANG (2015): “Self-Signaling and Pro-Social Behavior: a cause marketing mobile field experiment,” *Working Paper*.
- FONG, N. M., Z. FANG, AND X. LUO (2015): “Geo-Conquesting: Competitive Locational Targeting of Mobile Promotions,” *Journal of Marketing Research*, 52, 726–735.
- FUDENBERG, D., AND J. M. VILLAS-BOAS (2006): “Behavior-Based Price Discrimination and Customer Recognition,” in *Hanbooks in Information Systems: Economics and Information Systems*, ed. by T. Hendershott. Elsevier.

- GHOSE, A., A. GOLDFARB, AND S. P. HAN (2013): “How Is the Mobile Internet Different? Search Costs and Local Activities,” *Information Systems Research*, 24, 613–631.
- GOLDBERG, P. K., AND F. VERBOVEN (2005): “Market integration and convergence to the law of one price: Evidence from the European car market,” *Journal of International Economics*, 65, 49–73.
- HOLMES, T. J. (1989): “The Effects of Third-Degree Price Discrimination in Oligopoly Markets,” *The American Economic Review*, 79, 244–250.
- LUO, X., M. ANDREWS, Z. FANG, AND C. W. PHANG (2014): “Mobile Targeting,” *Management Science*, 60, 1738 – 1756.
- MCCULLOCH, R., AND P. E. ROSSI (1994): “An exact likelihood analysis of the multinomial probit model,” *Journal of Econometrics*, 64, 207–240.
- MIZUNO, T. (2003): “On the existence of a unique price equilibrium for models of product differentiation,” *International Journal of Industrial Organization*, 21, 761–793.
- NEVO, A., AND C. WOLFRAM (2002): “Why Do Manufacturers Issue Coupons? An Empirical Analysis of Breakfast Cereals,” *RAND Journal of Economics*, 33(2), 319–339.
- NEWTON, M. A., AND A. E. RAFTERY (1994): “Approximate Bayesian Inference with the Weighted Likelihood Bootstrap,” *Journal of the Royal Statistical Society, Series B*, 56(1), 3–48.
- PANCRAS, J., AND K. SUDHIR (2007): “Optimal Marketing Strategies for a Customer Data Intermediary,” *Journal of Marketing Research*, XLIV, 560–578.
- PIGOU, A. (1920): *The Economics of Welfare*. MacMillan.
- SHAFFER, G., AND Z. J. ZHANG (1995): “Competitive Coupon Targeting,” *Marketing Science*, 14, 395–416.
- SHEPARD, A. (1991): “Price Discrimination and Retail Configuration,” *Journal of Political Economy*, 99(1), 30–53.
- SHIN, J., AND K. SUDHIR (2010): “A Customer Management Dilemma : When is it Profitable to Reward Own Customers,” *Marketing Science*, 29, 671–689.
- STOLE, L. A. (2007): “Price Discrimination and Imperfect Competition,” in *Handbook of Industrial Organization*, vol. 3. Elsevier.
- THISSE, J.-F., AND X. VIVES (1988): “On the Strategic Choice of Spatial Price Policy,” *The American Economic Review*, 78, 122–137.
- VARIAN, H. R. (1989): *Price Discrimination* chap. 10, pp. 597–654. Elsevier Science Publishers BV.

Figure 1: Aggregate Purchase Rates

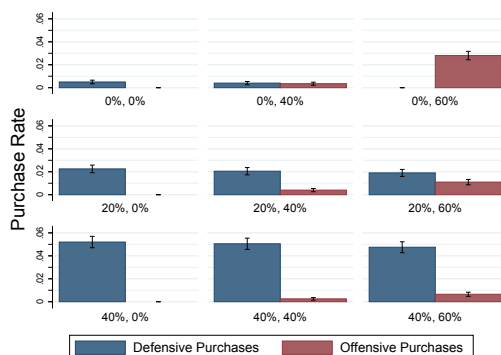


Figure 2: Expected revenues per consumer

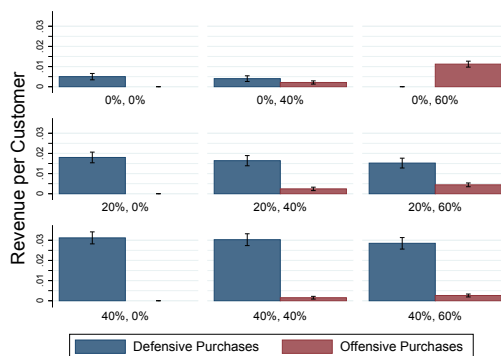


Table 1: Posterior Model Fit by Segment

segment	Multinomial Logit	Multinomial Probit
high consumers in location A	-778.5684	-774.7321
low consumers in location A	-456.5403	-456.4356
high consumers in location B	-784.3276	-768.8367
low consumers in location B	-489.8145	-488.7583

Figure 3: Purchase Rate by Segment



Figure 4: Expected revenues per consumer by segment



Table 2: Posterior Means of Multinomial Probit by segment (90% posterior credibility intervals in parentheses)

coefficient	high A	low A	high B	low B
θ_A	-0.344	0.25	-1.066	-1.413
	(-0.651,-0.028)	(-0.178,0.695)	(-1.344,-0.79)	(-1.737,-0.964)
θ_B	-1.043	-0.628	-0.376	0
	(-2.002,-0.425)	(-1.499,-0.023)	(-0.741,-0.035)	(-0.311,0.349)
α	-0.027	-0.044	-0.027	-0.028
	(-0.033,-0.021)	(-0.053,-0.035)	(-0.036,-0.019)	(-0.043,-0.017)
$\Psi_{A,A}$	1	1	1	1
	-	-	-	-
$\Psi_{B,B}$	1.006	0.738	1.152	0.577
	(0.437,2.105)	(0.323,1.393)	(0.692,1.651)	(0.287,1.237)
$\Psi_{A,B}$	0.787	-0.795	1.025	0.152
	(0.341,1.259)	(-1.125,-0.542)	(0.801,1.234)	(-1.019,0.663)
$\rho_{A,B}$	0.796	-0.951	0.962	0.348
	(0.443,0.931)	(-0.99,-0.826)	(0.926,0.985)	(-0.953,0.955)

Table 3: Multinomial Probit Elasticities by segment (evaluated at regular prices of 75RMB)

	high, A		low, A		high, B		low, B	
	both set regular prices of 75 RMB							
	Firm A price	Firm B price	Firm A price	Firm B price	Firm A price	Firm B price	Firm A price	Firm B price
Firm A	-5.33	0.15	-10.17	1.07E-16	-16.99	13.17	-7.88	3.72
Firm B	3.44	-8.35	1.77E-14	-11.82	0.02	-4.84	0.42	-8.96
	both set prices of 30 RMB (60% off)							
	Firm A price	Firm B price	Firm A price	Firm B price	Firm A price	Firm B price	Firm A price	Firm B price
Firm A	-1.40	0.10	-2.07	0.00	-7.97	5.95	-3.10	0.77
Firm B	1.52	-3.44	0.00	-4.33	0.01	-1.25	0.03	-1.91

Figure 5: Shift in Posterior Expected Demand for Theater A when Theater B cuts its price (dotted lines represent the 90% posterior credibility intervals)

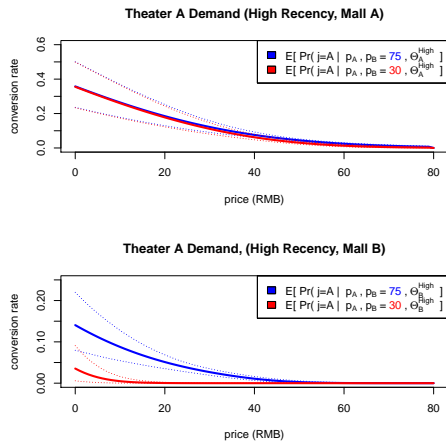


Figure 6: Shift in Posterior Expected Demand for Theater A when Theater B cuts its price
(dotted lines represent the 90% posterior credibility intervals)




Table 4: Uniform Price Equilibrium

		Firm A	Firm B
Price		19.2942	18.8641
Share:	high type, location A	0.1896	0.0168
	low type, location A	0.2795	0.0465
	high type, location B	0.0005	0.2039
	low type, location B	0.0106	0.2380
Expected Profit per messaged Customer		1.9604	2.9133

Table 5: Unilateral Targeting (90% posterior credibility intervals in parentheses)

	Firm A Profit per messaged customer	Firm B Profit per messaged customer
Uniform	1.96 (1.46,2.52)	2.91 (2.24,3.74)
Location	1.98 (1.47,2.55)	3.02 (2.34,3.85)
Type	1.97 (1.46,2.53)	2.94 (2.26,3.76)
Type and Location	2.00 (1.48,2.56)	3.04 (2.35,3.88)

Figure 7: Posterior Distribution of Percent Difference in Unilateral Revenues per Messaged Customer Under Targeting versus Uniform

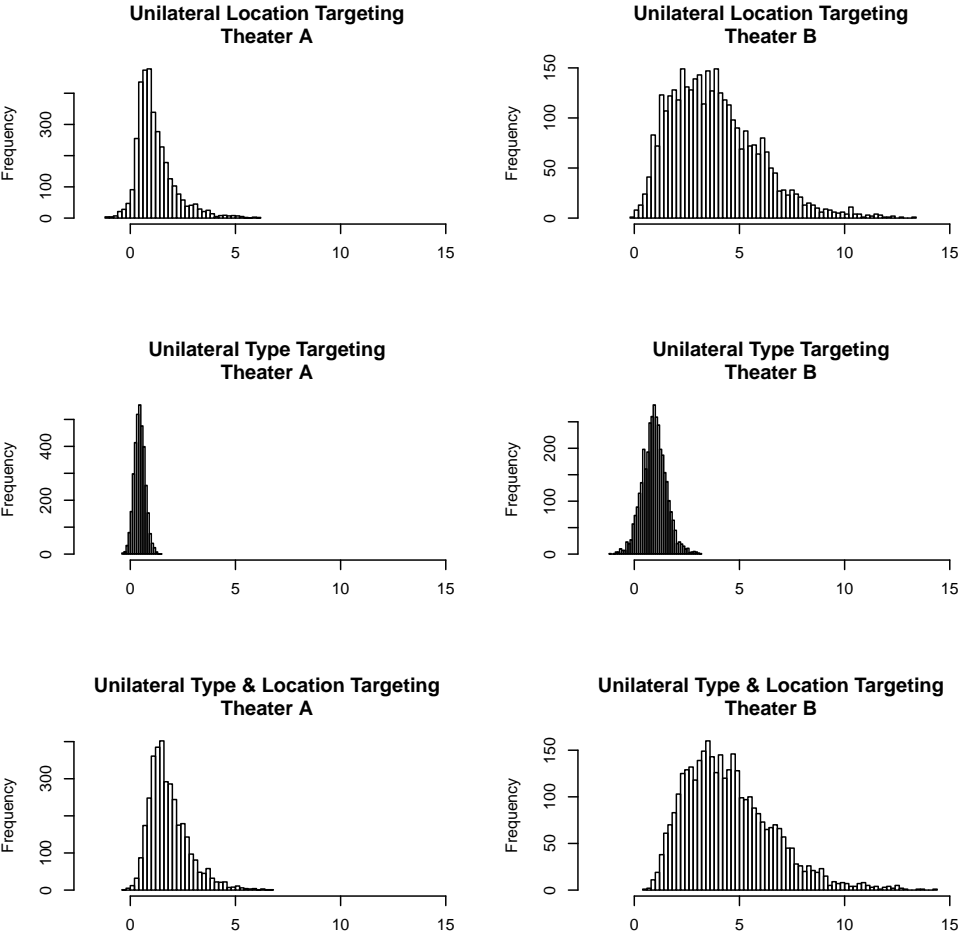


Table 6: Equilibrium Targeting (90% posterior credibility intervals in parentheses)

	Firm A Profit per messaged customer	Firm B Profit per messaged customer
Uniform	1.96 (1.46,2.52)	2.91 (2.24,3.74)
Location	1.96 (1.46,2.53)	2.98 (2.3,3.82)
Type	1.98 (1.47,2.54)	2.95 (2.27,3.77)
Type and Location	1.97 (1.47,2.54)	2.97 (2.28,3.8)

Figure 8: Best-Response Functions for Geographic Targeting

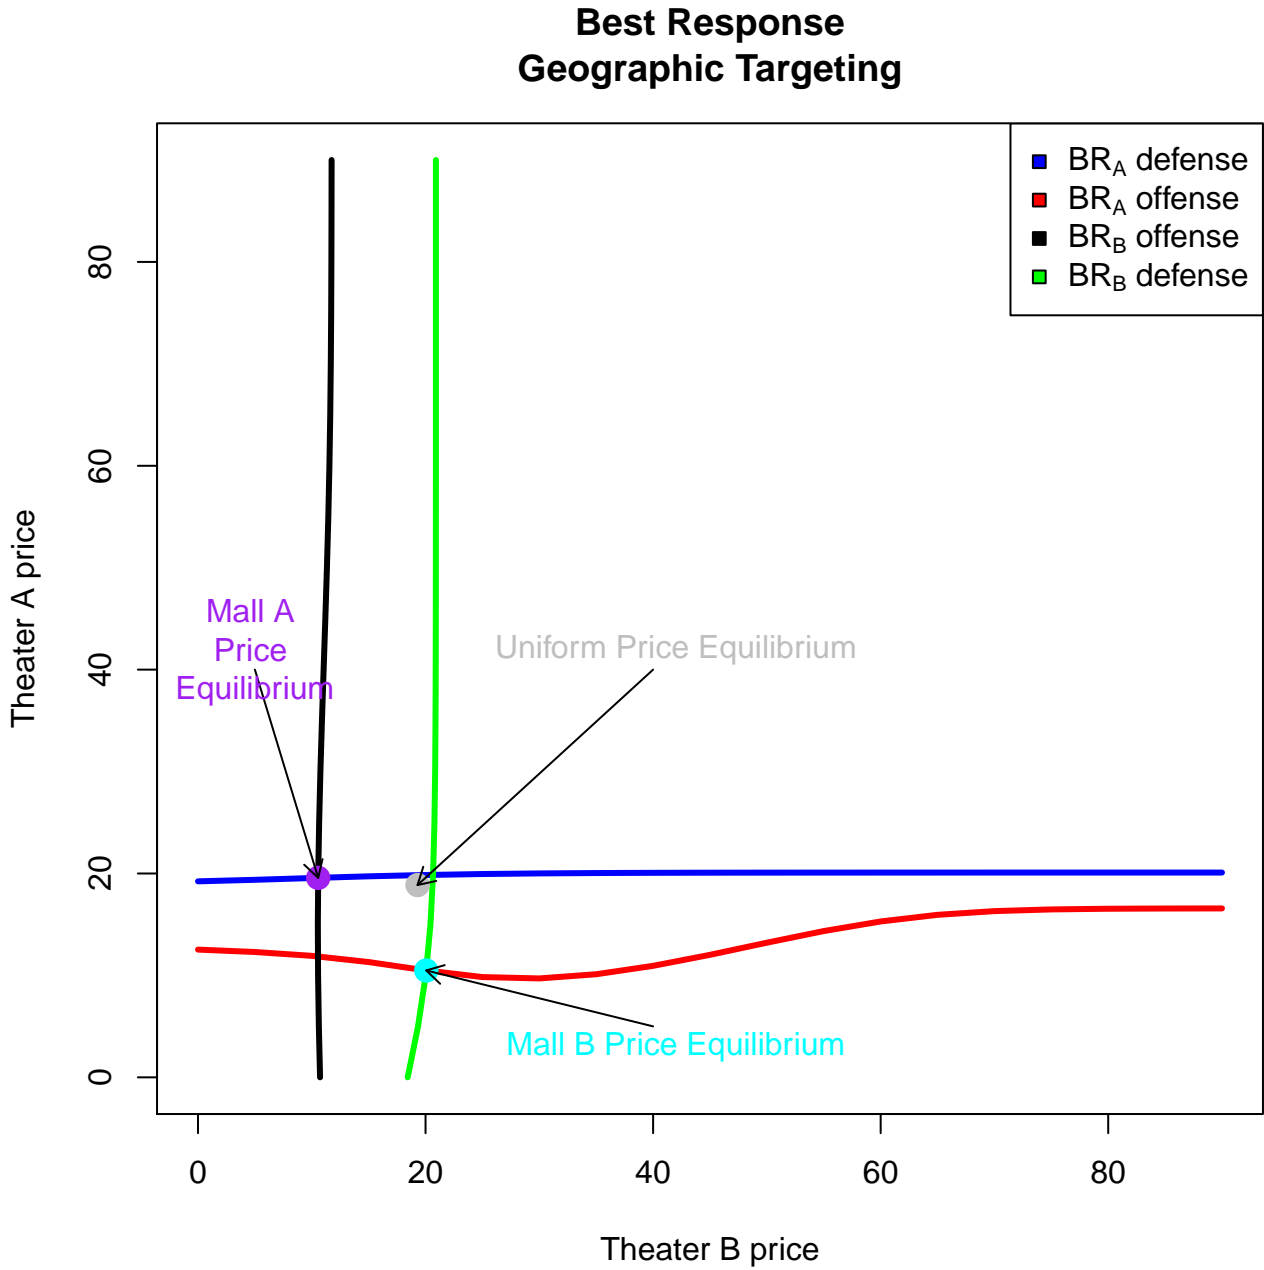


Figure 9: Best-Response Functions for Behavioral Targeting

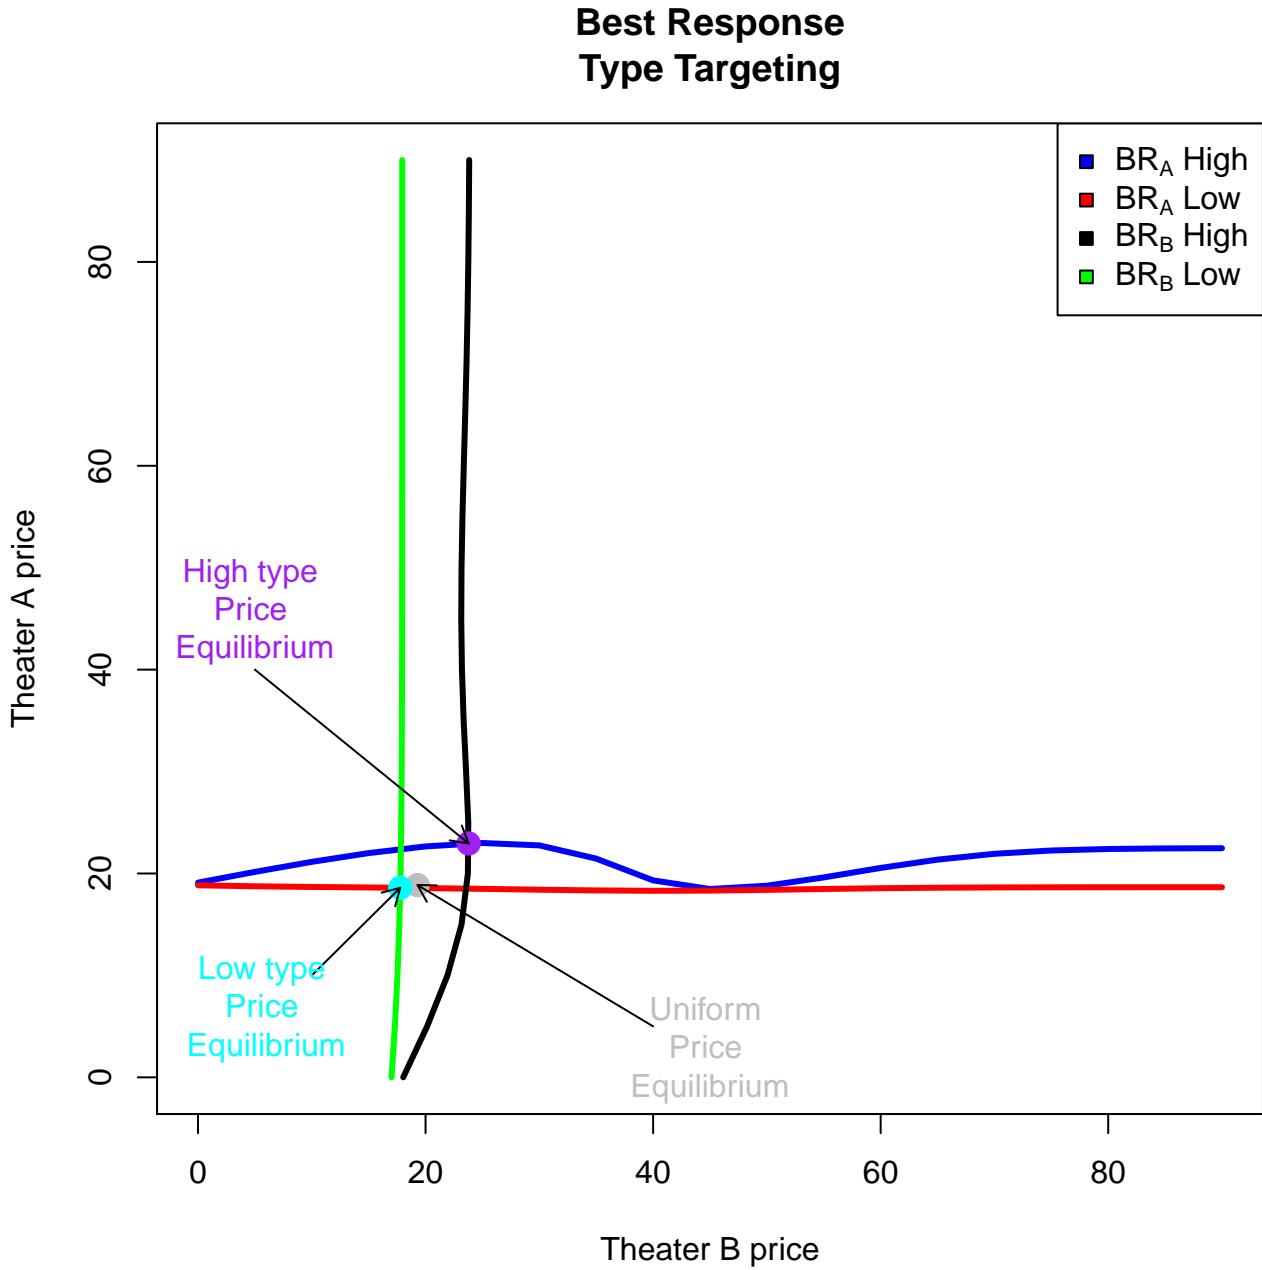


Table 7: Equilibrium Prices

	market	Firm A Price	Firm B Price
Uniform	pooled	19.294	18.864
by geography	Loc A	19.575	10.564
	Loc B	10.485	20.064
by type	High	22.948	23.786
	Low	18.597	17.775
by geography and type	Loc A, High	21.335	10.870
	Loc A, Low	19.146	10.546
	Loc B, High	5.230	20.595
	Loc B, Low	11.874	19.322

Figure 10: Posterior Distribution of Percent Difference in Equilibrium Revenues per Messaged Customer Under Targeting versus Uniform

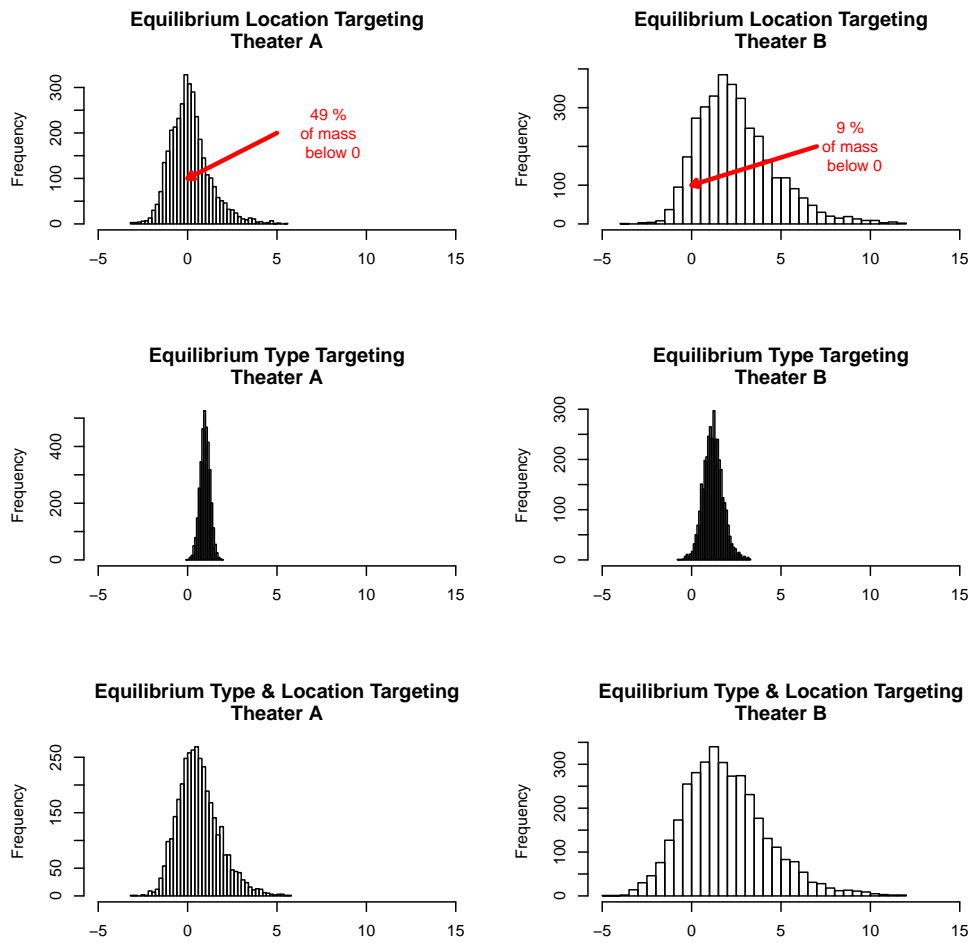


Table 8: Targeting Choice as a Strategic Game (firm payoffs per messaged customer)

(a) Targeting on Location					
		Firm B			
		Uniform		Targeted on type & location	
Firm A	Uniform	1.96	2.91	1.94	3.04
	Targeted on type & location	1.98	2.91	1.97	2.97
(b) Targeting on Type					
		Firm B			
		Uniform		Targeted on type	
Firm A	Uniform	1.96	2.91	1.97	2.94
	Targeted on type	1.97	2.92	1.98	2.95
(c) Targeting on Type and Location					
		Firm B			
		Uniform		Targeted on type	
Firm A	Uniform	1.96	2.91	1.94	3.01
	Targeted on type	2.00	2.80	1.96	2.98

A Appendix: Probit Derivatives

Recall that the expected probability that a consumer chooses alternative j is

$$Pr(y = j|B, X, \Psi) = \Phi\left(\mu_{z1}^{(1)}, \mu_{z2}^{(1)}; \rho_Z^{(1)}\right).$$

The matrix of derivatives of the share is as follows:

$$\frac{\partial Prob(y = j|B, X, \Psi)}{\partial X^T} = \frac{\partial Prob(y = j|B, X, \Psi)}{\partial \mu_z^{(j)T}} \frac{\partial \mu_z^{(j)}}{\partial X^T}$$

where

$$\frac{\partial \mu_z^{(j)}}{\partial X^T} = \text{diag}\left(\Sigma_z^{(j)}\right)^{-\frac{1}{2}} \Delta^{(j)} B.$$

It is straightforward to show (see below) that

$$\frac{\partial \Phi(x, y; \rho)}{\partial x} = \phi(x) \Phi\left(\frac{y - \rho x}{\sqrt{1 - \rho^2}}\right).$$

and therefore

$$\frac{\partial Pr(y = j|B, X, \Psi)}{\partial \mu_{zi}^{(j)}} = \phi\left(\mu_{zi}^{(j)}\right) \Phi\left(\frac{\mu_{z(3-i)}^{(j)} - \rho \mu_{zi}^{(j)}}{\sqrt{1 - \rho^2}}\right).$$

A.1 Derivative of bivariate Gaussian

$$\frac{\partial \Phi(x, y; \rho)}{\partial x} = \int_{-\infty}^y \frac{1}{2\pi\sqrt{1-\rho^2}} \exp\left(-\frac{x^2 - 2\rho xv + v^2}{2(1-\rho^2)}\right) dv.$$

If you complete the square inside the $\exp(\bullet)$ function, you can isolate the component depending on v :

$$\exp\left(-\frac{x^2 - 2\rho xv + v^2}{2(1-\rho^2)}\right) = \exp\left(-\frac{(v - \rho x)^2}{2(1-\rho^2)}\right) \exp\left(-\frac{x^2}{2}\right).$$

And so we can re-write the derivative as

$$\frac{\partial \Phi(x, y; \rho)}{\partial x} = \phi(x) \Phi\left(\frac{y - \rho x}{\sqrt{1 - \rho^2}}\right).$$

B Appendix: Supplemental Figures and Tables

Table 9: Summary Statistics. ARPU = “average revenue per user,” MOU = “average minutes used per month,” SMS= “average number of SMS sent per month,” GPRS= “average kilobytes downloaded per month.”

Segment	ARPU	MOU	SMS	GPRS	N
Loc A & High	109.96	771.54	205.23	90127.72	4450
	-85.19	-720.44	-279.56	-217276	
Loc A & Low	111.02	772.97	202	85707.8	4461
	-92.06	-726.6	-224.21	-132691	
Loc B & High	110.22	766.16	212.39	94731.77	4550
	-92.14	-709.3	-327.51	-274771.8	
Loc B & Low	112.19	774.46	206.72	90548	4539
	-87.47	-711.54	-271.03	-206697.4	
High	110.09	768.82	208.85	92455.32	9000
	-88.77	-714.82	-304.76	-248021	
Low	111.61	773.72	204.38	88148.87	9000
	-89.77	-719.03	-248.93	-174008.4	
Location A	110.49	772.26	203.61	87915.03	8911
	-88.69	-723.52	-253.37	-179981.3	
Location B	111.2	770.31	209.56	92642.42	9089
	-89.84	-710.42	-300.64	-243174.3	
all	110.85	771.27	206.61	90302.1	18000
	-89.27	-716.93	-278.26	-214243.91	

Table 10: Mobile Usage Randomization Checks

	ARPU	MOU	SMS	GPRS	Combined
Unadjusted P<.05	6	0	0	0	6
Adjusted P<.05	0	0	0	0	0
Number of Comparisons	36	36	36	36	144
Unadjusted Rate	17%	0%	0%	0%	4%
Adjusted Rate	0%	0%	0%	0%	0%

Note: Randomization checks for assignment of pricing treatments were performed using customers’ historical mobile usage variables presented in Table 9. Unadjusted P<.05 and Adjusted P<.05 count the number of pairwise comparisons between experimental cells where average mobile usages had statistically significant differences. The corresponding rates divides the counts by the number of comparisons. The unadjusted P-values find differences at an overall rate expected by chance. Adjusted P-values use Tukey’s honest significant difference adjustments for multiple comparisons of pairwise means; the adjusted P-values find no significant differences.

Table 11: Comparison of Locations

	Location A	Location B
Shopping Area (sq. meters)	102,000	120,000
Bus Lines	10	10
Visitors (people/day)	53,000	55,000
Number of Merchants	650	670
Population (1km radius)	26,367	24,233

Note: Shopping mall location statistics are drawn from the respective malls' promotional materials, except for population. The nearby population (within 1km) was estimated using GIS data from the 2010 Census, provided by a research center at the University of Michigan Ann Arbor.

Table 12: Aggregate Purchase Rates for Offensive Promotions

		Offensive Discount					
Defensive Discount		A 0%	B 40%	C 60%	D (B) - (A)	E (C) - (A)	F (C) - (B)
1	0%	0.0000 (0.0000)	0.0035** (0.0013)	0.0280** (0.0037)	0.0035** (0.0013)	0.0280** (0.0037)	0.0245** (0.0039)
2	20%	0.0000 (0.0000)	0.0040** (0.0014)	0.0110** (0.0023)	0.0040** (0.0014)	0.0110** (0.0023)	0.0070* (0.0027)
3	40%	0.0000 (0.0000)	0.0025* (0.0011)	0.0065** (0.0018)	0.0025* (0.0011)	0.0065** (0.0018)	0.0040† (0.0021)
4	(2) - (1)	0.0000 (0.0000)	0.0005 (0.0019)	-0.0170** (0.0044)	0.0005 (0.0019)	-0.0170** (0.0044)	-0.0175** (0.0048)
5	(3) - (1)	0.0000 (0.0000)	-0.0010 (0.0017)	-0.0215** (0.0041)	-0.0010 (0.0017)	-0.0215** (0.0041)	-0.0205** (0.0045)
6	(3) - (2)	0.0000 (0.0000)	-0.0015 (0.0018)	-0.0045 (0.0029)	-0.0015 (0.0018)	-0.0045 (0.0029)	-0.0030 (0.0035)

** p<.01, * p<.05, † p<.10, standard errors in parentheses

Note: Standard errors for differences in proportions and all p-values computed using conventional normal approximation. Since the approximation can perform poorly for very small proportions, we also test using several alternatives, including linear regression (conventional and robust standard errors), non-parametric bootstrap, and permutation testing, all of which obtain similar results (available from authors upon request). Sample size is 2,000 per cell (N=18,000 total).

Table 13: Aggregate Purchase Rates for Defensive Promotions

		Offensive Discount					
Defensive Discount		A 0%	B 40%	C 60%	D (B) - (A)	E (C) - (A)	F (C) - (B)
1	0%	0.0050** (0.0016)	0.0040** (0.0014)	0.0000 (0.0000)	-0.0010 (0.0021)	-0.0050** (0.0016)	-0.0040** (0.0014)
2	20%	0.0225** (0.0033)	0.0205** (0.0032)	0.0190** (0.0031)	-0.0020 (0.0046)	-0.0035 (0.0045)	-0.0015 (0.0044)
3	40%	0.0520** (0.0050)	0.0505** (0.0049)	0.0475** (0.0048)	-0.0015 (0.0070)	-0.0045 (0.0069)	-0.0030 (0.0068)
4	(2) - (1)	0.0175** (0.0037)	0.0165** (0.0035)	0.0190** (0.0031)	-0.0010 (0.0051)	0.0015 (0.0048)	0.0025 (0.0046)
5	(3) - (1)	0.0470** (0.0052)	0.0465** (0.0051)	0.0475** (0.0048)	-0.0005 (0.0073)	0.0005 (0.0071)	0.0010 (0.0070)
6	(3) - (2)	0.0295** (0.0060)	0.0300** (0.0058)	0.0285** (0.0057)	0.0005 (0.0083)	-0.0010 (0.0082)	-0.0015 (0.0081)

Figure 11: Offensive Purchase Rates vs. Logit and Probit Predicted Rates




Table 14: Purchase Rates for Location A High Types

Offensive Response							
		Offensive Discount					
Defensive Discount		A 0%	B 40%	C 60%	D (B) - (A)	E (C) - (A)	F (C) - (B)
1	0%	0.0000	0.0060†	0.0419**	0.0060†	0.0419**	0.0360**
		(0.0000)	(0.0034)	(0.0092)	(0.0034)	(0.0092)	(0.0098)
2	20%	0.0000	0.0084*	0.0156**	0.0084*	0.0156**	0.0072
		(0.0000)	(0.0042)	(0.0055)	(0.0042)	(0.0055)	(0.0069)
3	40%	0.0000	0.0040	0.0123*	0.0040	0.0123*	0.0083
		(0.0000)	(0.0028)	(0.0050)	(0.0028)	(0.0050)	(0.0057)
4	(2) - (1)	0.0000	0.0024	-0.0263*	0.0024	-0.0263*	-0.0288*
		(0.0000)	(0.0054)	(0.0107)	(0.0054)	(0.0107)	(0.0120)
5	(3) - (1)	0.0000	-0.0020	-0.0296**	-0.0020	-0.0296**	-0.0276*
		(0.0000)	(0.0044)	(0.0104)	(0.0044)	(0.0104)	(0.0113)
6	(3) - (2)	0.0000	-0.0044	-0.0033	-0.0044	-0.0033	0.0011
		(0.0000)	(0.0050)	(0.0074)	(0.0050)	(0.0074)	(0.0090)
Defensive Response							
		Offensive Discount					
Defensive Discount		A 0%	B 40%	C 60%	D (B) - (A)	E (C) - (A)	F (C) - (B)
1	0%	0.0123*	0.0100*	0.0000	-0.0023	-0.0123*	-0.0100*
		(0.0050)	(0.0044)	(0.0000)	(0.0067)	(0.0050)	(0.0044)
2	20%	0.0233**	0.0210**	0.0253**	-0.0023	0.0020	0.0043
		(0.0067)	(0.0066)	(0.0069)	(0.0094)	(0.0096)	(0.0096)
3	40%	0.0657**	0.0595**	0.0492**	-0.0062	-0.0165	-0.0103
		(0.0112)	(0.0105)	(0.0098)	(0.0154)	(0.0149)	(0.0144)
4	(2) - (1)	0.0111	0.0110	0.0253**	0.0000	0.0143	0.0143
		(0.0083)	(0.0079)	(0.0069)	(0.0115)	(0.0108)	(0.0105)
5	(3) - (1)	0.0534**	0.0496**	0.0492**	-0.0039	-0.0043	-0.0004
		(0.0123)	(0.0114)	(0.0098)	(0.0168)	(0.0157)	(0.0151)
6	(3) - (2)	0.0424**	0.0385**	0.0238*	-0.0038	-0.0185	-0.0147
		(0.0131)	(0.0124)	(0.0120)	(0.0180)	(0.0177)	(0.0173)

Table 15: Purchase Rates for Location A Low Types

Offensive Response							
		Offensive Discount					
Defensive Discount		A 0%	B 40%	C 60%	D (B) - (A)	E (C) - (A)	F (C) - (B)
1	0%	0.0000	0.0000	0.0185**	0.0000	0.0185**	0.0185**
		(0.0000)	(0.0000)	(0.0061)	(0.0000)	(0.0061)	(0.0061)
2	20%	0.0000	0.0000	0.0078*	0.0000	0.0078*	0.0078*
		(0.0000)	(0.0000)	(0.0039)	(0.0000)	(0.0039)	(0.0039)
3	40%	0.0000	0.0039	0.0041	0.0039	0.0041	0.0002
		(0.0000)	(0.0027)	(0.0029)	(0.0027)	(0.0029)	(0.0040)
4	(2) - (1)	0.0000	0.0000	-0.0107	0.0000	-0.0107	-0.0107
		(0.0000)	(0.0000)	(0.0072)	(0.0000)	(0.0072)	(0.0072)
5	(3) - (1)	0.0000	0.0039	-0.0144*	0.0039	-0.0144*	-0.0182*
		(0.0000)	(0.0027)	(0.0068)	(0.0027)	(0.0068)	(0.0073)
6	(3) - (2)	0.0000	0.0039	-0.0037	0.0039	-0.0037	-0.0075
		(0.0000)	(0.0027)	(0.0048)	(0.0027)	(0.0048)	(0.0056)
Defensive Response							
		Offensive Discount					
Defensive Discount		A 0%	B 40%	C 60%	D (B) - (A)	E (C) - (A)	F (C) - (B)
1	0%	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
		(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)
2	20%	0.0079*	0.0170**	0.0117*	0.0091	0.0037	-0.0053
		(0.0039)	(0.0060)	(0.0047)	(0.0071)	(0.0062)	(0.0076)
3	40%	0.0418**	0.0388**	0.0432**	-0.0030	0.0014	0.0044
		(0.0092)	(0.0085)	(0.0092)	(0.0125)	(0.0130)	(0.0126)
4	(2) - (1)	0.0079*	0.0170**	0.0117*	0.0091	0.0037	-0.0053
		(0.0039)	(0.0060)	(0.0047)	(0.0071)	(0.0062)	(0.0076)
5	(3) - (1)	0.0418**	0.0388**	0.0432**	-0.0030	0.0014	0.0044
		(0.0092)	(0.0085)	(0.0092)	(0.0125)	(0.0130)	(0.0126)
6	(3) - (2)	0.0339**	0.0219*	0.0316**	-0.0121	-0.0024	0.0097
		(0.0100)	(0.0104)	(0.0104)	(0.0144)	(0.0144)	(0.0147)

Table 16: Purchase Rates for Location B High Types

Offensive Response							
		Offensive Discount					
Defensive Discount		A 0%	B 40%	C 60%	D (B) - (A)	E (C) - (A)	F (C) - (B)
1	0%	0.0000 (0.0000)	0.0080* (0.0040)	0.0363** (0.0082)	0.0080* (0.0040)	0.0363** (0.0082)	0.0283** (0.0091)
2	20%	0.0000 (0.0000)	0.0000 (0.0000)	0.0144** (0.0054)	0.0000 (0.0000)	0.0144** (0.0054)	0.0144** (0.0054)
3	40%	0.0000 (0.0000)	0.0000 (0.0000)	0.0039 (0.0028)	0.0000 (0.0000)	0.0039 (0.0028)	0.0039 (0.0028)
4	(2) - (1)	0.0000 (0.0000)	-0.0080* (0.0040)	-0.0220* (0.0098)	-0.0080* (0.0040)	-0.0220* (0.0098)	-0.0139 (0.0106)
5	(3) - (1)	0.0000 (0.0000)	-0.0080* (0.0040)	-0.0324** (0.0086)	-0.0080* (0.0040)	-0.0324** (0.0086)	-0.0244* (0.0095)
6	(3) - (2)	0.0000 (0.0000)	0.0000 (0.0000)	-0.0105† (0.0061)	0.0000 (0.0000)	-0.0105† (0.0061)	-0.0105† (0.0061)
Defensive Response							
		Offensive Discount					
Defensive Discount		A 0%	B 40%	C 60%	D (B) - (A)	E (C) - (A)	F (C) - (B)
1	0%	0.0078* (0.0039)	0.0060† (0.0035)	0.0000 (0.0000)	-0.0018 (0.0052)	-0.0078* (0.0039)	-0.0060† (0.0035)
2	20%	0.0370** (0.0086)	0.0344** (0.0080)	0.0287** (0.0076)	-0.0027 (0.0117)	-0.0083 (0.0114)	-0.0056 (0.0110)
3	40%	0.0585** (0.0104)	0.0605** (0.0107)	0.0605** (0.0105)	0.0020 (0.0149)	0.0021 (0.0148)	0.0001 (0.0150)
4	(2) - (1)	0.0292** (0.0094)	0.0283** (0.0087)	0.0287** (0.0076)	-0.0009 (0.0128)	-0.0005 (0.0121)	0.0004 (0.0115)
5	(3) - (1)	0.0507** (0.0111)	0.0545** (0.0113)	0.0605** (0.0105)	0.0038 (0.0158)	0.0099 (0.0153)	0.0061 (0.0154)
6	(3) - (2)	0.0214 (0.0134)	0.0261† (0.0133)	0.0318* (0.0130)	0.0047 (0.0189)	0.0104 (0.0187)	0.0057 (0.0186)

Table 17: Purchase Rates for Location B Low Types

Offensive Response							
		Offensive Discount					
Defensive Discount		A 0%	B 40%	C 60%	D (B) - (A)	E (C) - (A)	F (C) - (B)
1	0%	0.0000	0.0000	0.0156**	0.0000	0.0156**	0.0156**
		(0.0000)	(0.0000)	(0.0055)	(0.0000)	(0.0055)	(0.0055)
2	20%	0.0000	0.0076*	0.0062†	0.0076*	0.0062†	-0.0014
		(0.0000)	(0.0038)	(0.0036)	(0.0038)	(0.0036)	(0.0052)
3	40%	0.0000	0.0021	0.0058†	0.0021	0.0058†	0.0038
		(0.0000)	(0.0021)	(0.0034)	(0.0021)	(0.0034)	(0.0039)
4	(2) - (1)	0.0000	0.0076*	-0.0094	0.0076*	-0.0094	-0.0170*
		(0.0000)	(0.0038)	(0.0065)	(0.0038)	(0.0065)	(0.0075)
5	(3) - (1)	0.0000	0.0021	-0.0098	0.0021	-0.0098	-0.0118†
		(0.0000)	(0.0021)	(0.0064)	(0.0021)	(0.0064)	(0.0067)
6	(3) - (2)	0.0000	-0.0055	-0.0003	-0.0055	-0.0003	0.0052
		(0.0000)	(0.0043)	(0.0049)	(0.0043)	(0.0049)	(0.0065)
Defensive Response							
		Offensive Discount					
Defensive Discount		A 0%	B 40%	C 60%	D (B) - (A)	E (C) - (A)	F (C) - (B)
1	0%	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
		(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)
2	20%	0.0223**	0.0095*	0.0103*	-0.0128	-0.0120	0.0009
		(0.0066)	(0.0042)	(0.0046)	(0.0079)	(0.0081)	(0.0062)
3	40%	0.0421**	0.0433**	0.0370**	0.0012	-0.0052	-0.0063
		(0.0088)	(0.0092)	(0.0083)	(0.0128)	(0.0121)	(0.0124)
4	(2) - (1)	0.0223**	0.0095*	0.0103*	-0.0128	-0.0120	0.0009
		(0.0066)	(0.0042)	(0.0046)	(0.0079)	(0.0081)	(0.0062)
5	(3) - (1)	0.0421**	0.0433**	0.0370**	0.0012	-0.0052	-0.0063
		(0.0088)	(0.0092)	(0.0083)	(0.0128)	(0.0121)	(0.0124)
6	(3) - (2)	0.0199†	0.0338**	0.0267**	0.0140	0.0068	-0.0072
		(0.0110)	(0.0102)	(0.0095)	(0.0150)	(0.0145)	(0.0139)

Table 18: Equilibrium Targeting with $\rho = 0$

	Firm A Profit per messaged customer	Firm B Profit per messaged customer
Uniform	2.22	2.93
Location	2.19	3.00
Type	2.23	2.96
Type and Location	2.20	3.03

Figure 12: Defensive Purchase Rates vs. Logit and Probit Predicted Rates

