Title
Essays on Social Promotions and E-commerce

Author
Chen, Xin

Publication Date
2020

Peer reviewed|Thesis/dissertation

Essays on Social Promotions and E-commerce

by

Xin Chen

A dissertation submitted in partial satisfaction of the
requirements for the degree of
Doctor of Philosophy

in

Business Administration

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Ganesh Iyer, Co-chair
Professor Steve Tadelis, Co-chair
Associate Professor Zsolt Katona
Associate Professor Ben Handel

Spring 2020

Essays on Social Promotions and E-commerce

Abstract

Essays on Social Promotions and E-commerce

by

Xin Chen

Doctor of Philosophy in Business Administration

University of California, Berkeley

Professor Ganesh Iyer, Co-chair

Professor Steve Tadelis, Co-chair

Chapter one examines the effects of bundling social incentives with promotions. When using these "social promotions," a firm must decide whether or not to require customers to share the promotion on social media with friends. On the one hand, this requirement may make it easier for the firm to get the word out. On the other hand, this requirement may inconvenience customers and lower their propensity to purchase. To conduct valid inference on this trade-off between the costs and benefits, I devise novel empirical strategies that flexibly accommodate the amount of information available to a firm. Through a field experiment on a Chinese online grocery site, randomized both at the city and individual level, I find that social promotions can serve as an effective growth strategy. The social obligation can benefit a firm through two channels. First, the customers who do share add additional value through both new customer acquisition and existing customer retention. Second, the firm does not lose profit from the customers who choose not to purchase the promoted products when forced to share (but would otherwise purchase the promoted products without mandatory sharing); they contribute instead by purchasing other substitutes from the same site. Furthermore, different types of customers respond to the social incentive differently and I present results on customer heterogeneity.

Chapter two expands on the targeting question in social promotions. This paper examines how a firm should target a social incentive on top of a standard promotion, i.e. also requiring the customer to share the promotion with friends on social media to qualify for the discounts. A major challenge in this targeting question is that the targeting decision on a focal customer also affects how many leads will be generated. Therefore, tackling this targeting question requires a different framework compared with that of standard targeting. This paper compare a few targeting schemes, including a post-experiment strategy based on an "effective" profit idea and two contextual bandits that balance exploration and exploitation real time, namely LinUCB and LinTS. We find that in our offline simulations, the two contextual bandits perform similarly and better than post-experiment targeting. Targeting the social incentive to certain type of customers can significantly increase firm profitability.

Chapter three discusses an empirical case study on bundle design and pricing. Unlike what conventional bundle pricing theory suggests, we find that on the largest Chinese e-commerce platform Taobao.com, a bundle tends to cost more than purchasing each bundle item from the same online market separately. These bundle premiums are prevalent in multiple product categories, such as digital cameras, iPads, cell-phones, and video game consoles. Conservatively

speaking, in the digital single-lens reflex cameras market, "premiumed bundles" account for 30% of the market share. Furthermore, the magnitude of the premium is also beyond that of rational expectation. For example, for Canon 700D bundles, the average premium is 17 USD, or 4% the base good price. These bundle premiums can have significant implications for consumer welfare and platform design. This paper systematically documents these empirical regularities.

To my parents, my advisors, and my research team

Acknowledgments

I owe my deepest gratitude to my parents and my advisors for their continued support throughout my transition from a mathematician to an economist and marketer. I am grateful for Professor Ganesh Iyer for helpful early-stage discussions for my main project, an adventure in its own end, and also his constant support at all stages. I thank Professor Steve Tadelis for providing me with the opportunity to experience the e-commerce sector. The experience inspired me to establish my career in this stimulating industry and helped me develop the skills to work with a real business for my dissertation. I thank Professor Zsolt Katona and Ben Handel for their helpful coaching on improving my presentation skills. In addition to my committee members, I am also grateful for other Berkeley-Haas professors, especially Professor John Morgan. I acknowledge fellow PhD students (especially my co-author Zachary Zemin Zhong for teaching me the nitty-gritty of running regressions), extended family, and friends for helpful discussions. I thank Yunhao Huang, Peiyao Li, Yi Liu, Suyang Lu, Alicia Ting Luo, and Frank Jiazhong Mei for their excellent research assistance. I also acknowledge the Berkeley-Haas Fisher Center of Business Analytics for financial support. Finally, I thank the management and staff of the cooperating firm for their efforts during the implementation of this study. There was no financial conflict of interest in the implementation of the study; no author was compensated by the partner company or by any other entity for the production of the research. All errors are my own.

Contents

1

Social Promotional Strategies and Consumer Choice: Evidence from an Online Field Experiment

1.1 Introduction

Online social interactions have become ubiquitous and marketers are actively trying to leverage these connections among customers for business growth. Social advertising serves as a critical new customer acquisition channel. Indeed, 93% of marketers are currently using influencer marketing and 60% of marketers said they plan to increase their influencer marketing budgets in 2019 [Insights, 2018, Marketer, 2019]. Many firms have been experimenting with innovative social strategies to grow and activate their customer bases. One noticeable example is the leading Chinese social e-commerce firm Pinduoduo (PDD). Pinduoduo lets people shop together online and earn group discounts, and was valued at $1.5 billion USD within 21 months of launching [Elstrom and Ramli, 2017]. The practice of bundling promotions with social incentives is used in markets ranging from online retail to financial and educational services. Despite the ubiquity of these practices, no prior literature discusses why and to what extent are these social promotions successful.

Firms that use social promotions face several important design decisions. In particular, they have to decide whether or not to force consumers to share information with their friends on social media (hereafter, I refer to this sharing requirement as the "social incentive"). On the one hand, this may make it easier for the firm to get the word out about its products and offers. The firm will also accumulate additional assets, such as network data on customers' social interactions. On the other hand, the sharing requirement may inconvenience customers and lower their propensity to purchase from the firm. In addition to these direct effects, the social incentive may create peer effects,[1] in that a customer who is forced to share might respond differently if a different fraction of peers are required to share. Given the potential costs and benefits discussed above (with more background discussion to follow in section 1.2), when deploying social promotions, a firm needs to choose the fraction, or if possible, the segments of customers upon which to impose the social incentive.

[1]Also referred to as endogenous effects in Manski [2013], indirect effects, general equilibrium effects, externalities, interaction effects, or peer effects on the non-treated.

This paper provides both novel methods and substantive insights into the trade-off between the costs and benefits of social promotions. Despite the increasing ease of carrying out field experiments to overcome endogeneity issues, several challenges persist in providing causal answers to the questions of interest herein. This paper addresses the challenges and contributes novel empirical strategies that flexibly accommodate the amount of information available to a firm. It reports results from a two-level controlled experiment executed online across 65 cities in China (randomizing both at the city and individual levels). The experiment shows that when a firm engages in promotions, bundling them with social incentives can be beneficial, and the benefits can come from two distinct channels. The experiment also sheds light on what types of customers can be less averse to sharing and more influential in generating leads.

I propose an empirical strategy that causally evaluates the social impact without requiring access to additional data on customer characteristics or network information. This is especially useful for start-ups that wish to utilize social promotions to acquire customers in new markets they wish to enter, because they are frequently confronted with the cold-start problem of not having enough information on their customers. To conduct valid inference, the proposed strategy combines a two-level randomization design with a parsimonious model.

The challenges of measuring the social incentive's overall impact are severe. Standard experimental design and analysis approach do not directly apply. First, for the design, neither market-level randomization nor individual-level randomization alone suffices. First, the issue of applying market-level experimentation with a standard diff-in-diff research design is that this is likely under-powered. A firm may not operate in a large enough number of mutually independent markets and there can also be substantial heterogeneity across these markets. Despite the best attempt to improve the precision of the estimates by properly stratifying these markets based on the observed market characteristics, one is still likely to suffer from a low signal-to-noise ratio. Second, the potential presence of peer effects makes mere individual-level randomization inadequate. As a reminder, peer effects refer to the phenomenon that a customer who is forced to share might respond differently if a different fraction of peers are required to share. The assumption that there are no peer effects (aka. the Stable Unit Treatment Value Assumption or SUTVA), is usually made implicitly in other experimentation settings. It, however, may be violated in the context of social promotions. When the relative magnitude of peer effects is large enough (compared with direct treatment effects), inattention to the peer effects can result in a sub-optimal social promotional strategy. Therefore, measuring the overall social impact requires a two-level design that randomizes both at the city and individual level.[2]

Moreover, unlike other contexts in which multi-level randomization design is the state-of-the-art, understanding the social efficacy also requires a formal model. First, distinct from other contexts of interest, typically in policy analysis, customers in my setting are not given as a fixed pool before the intervention. Instead, they arrive dynamically over time and the traffic size itself may also depend on the market-level promotional intensity (the fraction of customers required

[2] To the best of my knowledge, this is the first paper that introduces a multi-level randomization design to a marketing strategy context. The two-level randomization design does find widespread applications in other fields, such as labor economics and political science. For examples, Duflo and Saez [2003] use a multi-level design to study spillovers in individuals' choice of retirement plan; Hudgens and Halloran [2008] adopt a two-level design to look at housing vouchers and vaccines; Miguel and Kremer [2004] study spillovers in the treatment of worms; Sinclair et al. [2012] study interpersonal communication in the context of voter-mobilization. There is also a growing literature on the methodological front-end examining how to properly address interference assumptions in network experiments (see Eckles et al. [2017] for a recent example).

to share in a market). Second, the social channel makes the outcome of interest a customer's "effective" contribution, including that from both herself and her referrals. However, standard regression analysis on this outcome would double count the referrals' contribution. Thus, to conduct valid inference, it is necessary to specify a formal model.

To demonstrate how my proposed strategy applies to the research questions of interest, I ran a field experiment in the context of online grocery shopping in China. The corporate partner is an emerging "unicorn" that targets households with income well above the national average. It runs social promotions on its signature high-velocity[3] products, which are mostly fashionable imported exotic fruits. The promoted products used in this study are Thai young coconuts. The primary social channel utilized is the popular social network WeChat. To enjoy the social promotion, a customer needs to log in to the grocery site and then share the coupon on WeChat (either on the news feed or within a group or directly with a friend). Once the customer fulfills the social incentive, the coupon is added to her account and ready for redemption. She will also be prompted to navigate back to the grocery site (see Appendix Figure A.1).

This experiment randomizes both at the city and individual level. It allows for identifications of both direct treatment effects and indirect peer effects. First, I assume that each city is an independent market and randomly assign the 65 cities (to which the firm delivers coconuts) into control and treatment groups. For a city in the control group, all customers are required to share. For a city in the treatment group, a random half of the customers are required to share. To improve the precision of estimates for the average treatment effects of interest, I divide the cities into two strata[4] by the first-tier criterion[5]: four first-tier cities (Beijing, Shanghai, Guangzhou, and Shenzhen) and the remaining 61 non-first-tier cities. Then I randomly assign each city to control or treatment within each stratum. Within a treated city, I randomly assign half of the customers to control and the other half to treatment. While a control customer has to share in order to qualify for the deal, a treated customer can directly enjoy the discount without sharing (see Figure A.2 for illustration of individual-level randomization in a treated city). For brevity, I refer to a control or a treated customer as a Sharing-required (S) customer or a No-Share (NS) customer, respectively. Regardless of whether a customer comes to the coupon page directly via the home page or through a friend's referral in WeChat, she will be randomized.[6] Based on the design, aggregate-level analysis compares the effectiveness of forcing all versus a random half of the customers to share.

Summary of Findings

Social promotions can be an effective growth strategy because the social incentive can benefit a firm through two channels. First, the customers who do share add additional value through

[3]A "high-velocity" product (or "baokuan" in Chinese) is an item that tends to sell out quickly and has high sales. It is a commonly encountered term on Chinese e-commerce sites. This type of product is usually displayed in a format similar to that of a best-seller on western e-commerce sites.

[4]In statistics, a stratum (plural strata) refers to a subset of the population which is being sampled.

[5]According to Wikipedia, the Chinese city tier system is a hierarchical classification of Chinese cities based on differences in consumer behavior, income level, population size, consumer sophistication, infrastructure, talent pool, and business opportunity.

[6]Given that a customer who wishes to redeem the coupon needs to sign in, I am able to make sure that the same customer is always assigned to the same treatment or control condition. Before a customer logs in, she simply sees "click to get coupon" on the coupon page and she is not assigned to treatment or control until she logs in (see Appendix Figure A.2 for the UI flow a customer sees).

both new customer acquisition and existing customer retention. Second, the firm does not lose profit from the customers who choose not to purchase the promoted products when forced to share (but would otherwise purchase the promoted products without mandatory sharing); they contribute instead by purchasing other substitutes from the same site. As expected, requiring sharing reduces a customer's propensity to purchase the promoted products by 21%. However, the requirement does not significantly change a customer's propensity to make a purchase from the platform (and the corresponding expenditure and profit). When forced to share, some customers substitute similar products within the same product category from the platform. This suggests that the social incentive can induce purchases of non-promoted products.[7] These results echo findings in Sahni et al. [2016], who showed that targeted coupon offers via email can serve as a form of advertising and can result in spillover purchases of non-promoted products. To summarize, for a customer considering the deal, removing the social incentive does not significantly change her own propensity to purchase (including both promoted and non-promoted products) and yet it lowers her lead-generation potential.

On the aggregate level, using the design and framework I devise, I find that the firm is better off forcing all versus only a random half of the customers to share. After adjusting for referrals' purchases, each customer's "effective" propensity to buy any products from the platform decreases by 4% if the firm randomly removes her social incentive with a 50% probability. The relative magnitude measured in terms of "effective" expenditure per customer is also 4%.

Moreover, inattention to the peer effects (that a customer who is forced to share might respond differently if a different fraction of her peers are required to share) may result in a suboptimal decision. For example, a firm that plans to maximize sales of the promoted products would choose the wrong strategy if the relative magnitude of the peer effect on the propensity to buy were 30% higher than the direct treatment effect. Second, I find heterogeneity in customers' propensity to share and lead-generation potential. At the city level, while requiring sharing helps with new customer acquisition and existing customer retention in both first-tier and non-first-tier cities, the contribution are significantly higher in first-tier than non-first-tier cities. This may suggest that consumer purchasing power and the breadth of the network may matter more than the strengths of the social ties. At the individual level, constructing characteristics based on customers' preferences and behavior, I find that the following types of customers tend to have higher propensity to share and lead-generation potential: those who are more price sensitive, who have been registered for a longer period of time, who have consumed the promoted products before, who have used gift cards less in the past, who have made a recent purchase longer ago, and who have purchased less frequently in the past year.

The rest of the paper proceeds as follows. Section 1.2 expands on the background for the different treatment effects of interest. Section 1.3 discusses the related literature. Section 1.4 elaborates on the institutional context. Section 1.5 presents balance tests, defines all notations used in later discussion, and shows the benefits of the two-level randomization. Section 1.6 presents identification strategies and reports average treatment effects on both the intensive and extensive margins (regarding consumer choice on the promoted products). Section 1.7 presents the proposed strategy to evaluate the social incentive's overall impact and reports results related to the promoted products. Subsection 1.7 of this section illustrates how inattention to peer

[7]Share aversion is the main reason behind why some customers purchase only non-promoted products. Among S customers who only purchase non-promoted products from the platform (within treated cities), only about 30% (104 out of 337 customers) choose to share. To put this in perspective, among all S customers (within 50% NS treated cities), only 5% choose to share and yet do not purchase the promoted products.

effects may result in a sub-optimal strategy of social promotions. Section 1.8 reports the overall social impact on all products and discusses the associated mechanisms thereof. Section 1.9 presents the empirical strategy and findings related to customer heterogeneity. Section 1.10 concludes.

1.2 Background

In conceptual or theoretical terms, the trade-off between the benefits and costs in social promotions can be understood with respect to two margins: the intensive and the extensive margins. On the one hand, the sharing requirement may negatively (or positively) affect a focal customer's propensity to share, and subsequent purchase decision, expenditure, and profit contribution. I refer to these within-customer outcomes as the intensive margins. On the other hand, if a customer does decide to share, the leads generated through the social channel may compensate for any potential profit loss due to mandatory sharing. I refer to this across-customer aspect as the extensive margin.

Intensive Margins

First, I discuss how a customer may respond to the social incentive. It is unclear whether a customer will increase or decrease her consumption without versus with the sharing requirement.

On the one hand, the social incentive may reduce a customer's consumption. This is because the sharing requirement may create disutility for customers for a variety of reasons. First, the desire for exclusivity can make consumers seek goods that distinguish themselves from the masses. Therefore, they may not want their friends to consume similar goods. Second, a customer may incur a hassle cost for fulfilling the sharing requirement; she has to go to social media, share and then return to the shopping site to make purchases. Thirdly, reputational concerns could make customers unwilling to share. For example, quality uncertainty about the platform and its products could make consumers hesitant to share. Consumers may also be afraid of spamming friends' news feed or inbox with irrelevant information.[8] This disutility may result in a lower coupon redemption rate and loss of sales of the promoted products. Thus, the social incentive will reduce firms' profitability, unless a customer decides to buy the promoted products without the coupon or substitute and purchase more profitable products from the firm. However, I have verified that among customers exposed to the deal, negligibly few buy the promoted products at the original price.

On the other hand, the sharing requirement might also directly increase customers' intent to purchase. One possible reason is the standard rationale described in Milgrom and Roberts [1982] and Bagwell and Riordan [1991]; a customer may interpret the social incentive as a signal of high product quality. When a firm decides on which product to use for promotion, it should pick a high-quality one. If the firm instead uses a low-quality product, then the informed customers will have low propensity to share and low purchase intent. Understanding this, uninformed customers can infer that a firm will only use a high-quality product in social

[8] In fact, many firms are well aware of consumers' potential aversion to sharing; some customers delete the shared item just seconds after fulfilling the sharing requirement. To prevent a customer from immediately deleting the shared ad after redeeming the promised benefits, some firms even mandate that the shared item needs to stay in the news feed for at least three days or a week and require another news feed screenshot for further verification.

promotions. Another possible reason is that the sharing requirement may help a customer derive a warm glow. The reason for this is that by sharing, a customer also helps her friends reduce search costs, consume products that match their tastes, and enjoy good deals. Finally, she might also enjoy utilizing this opportunity to show off her own good taste and social status; in my specific context, showing off a healthy and fashionable diet.

Extensive Margins

I now discuss the other aspect of the trade-off, i.e. the extensive margins. The sharing requirement allows a firm to leverage existing customers' social capital to reach other customers. I define the primary measure in this family of outcomes as the number of children customers a focal customer produces when she is forced to share (hereafter, I use the terms "child customer," "lead," and "referral" interchangeably). Other similar outcomes of interest include the number of new, existing, and buying customers introduced to a product. Customers' information sharing behavior may affect peers' decision making. It can inform peers about the existence or the quality of a product.

Moreover, those who are familiar with WeChat, the platform used for sharing, may be aware that there is a "private" mode that suffices for fulfilling the social incentive; a customer may choose to make the shared content visible to only herself (see Appendix Figure A.3). This institutional feature itself makes identifying extensive margins practically relevant.

Peer Effects

Furthermore, the same customer who is forced to share might respond differently if a different fraction of peers are required to share. The promotional intensity may affect the frequency at which a customer sees a shared coupon on WeChat and therefore her utility. The intensity can affect both the intensive and extensive margins in either direction. Inattention to the peer effects can result in sub-optimal managerial decisions.

Following the partial population experiments literature [Duflo and Saez, 2003], I define peer effects as differences between groups facing different probabilities of treatment. More precisely, the effects of interest are the average differences in outcomes between untreated units in treated groups and untreated units in pure control groups. A control customer is an S customer who is forced to share and a treated customer is an NS (No-Share) customer who receives a regular coupon without the social incentive. A treated group is a city where a random half of the customers are assigned to NS and a pure control group is a city where all customers are assigned to S. Therefore, the peer effects of interest here are the difference in outcomes between the S customers in 50% NS cities and S customers in the 0% NS cities.

Peer effects can affect the intensive margins in both directions. On the one hand, imposing social incentives on a larger fraction of customers can increase a customer's purchase intent and consumption amount. This can happen through at least three channels. First, there might be a pure advertising intensity effect. The more frequently a customer sees a shared coupon, the more likely she is to purchase and spend more. Second, there could also be a stronger-tie effect [Manchanda et al., 2015]. The more shared coupons a customer sees, the more likely she will be to see one from a close friend. If a stronger connection correlates with greater social influence, then she is more likely to use such a shared coupon. Third, given that more peers are required to share, potential "peer-induced" unfairness (Nguyen and Simkin, 2013, Li and

Jain, 2015) occurs with lower probability. This is because, an S customer may be dissatisfied after finding out that the same benefit/incentive is offered to a friend who does not incur any social incentive. On the other hand, imposing social incentives on a larger fraction of customers can also decrease a customer's purchase intent and consumption amount. She might want to assume her own social identity and stay unique by not conforming to friends' preferences.

For similar reasons, peer effects can also affect the extensive margins in both directions. An additional relevant factor here is the network structure. When customers' networks overlap substantially, a higher intensity may also decrease the average number of leads each customer generates.

Identification of the peer effects is critical for properly measuring a social promotional strategy's overall social impact. While Burszytn et al. [2014] show that peer effects can manifest themselves through different channels in the context of high-risk financial investment, the extent to which peer effects can affect firms' marketing strategies in the context of consumer packaged goods consumption remains to be better understood. Although I do not detect statistical significance on peer effects within my one-day experiment, this could be a false negative.

1.3 Related Literature

I contribute to the extant literature both substantively and methodologically. Substantively, I provide insights into how requiring information sharing in social promotions affects consumer behavior and firm strategies. Methodologically, to causally evaluate the social incentive's overall impact, I devise novel empirical strategies that flexibly accommodate the amount of information available to a firm.

The paper enriches the word-of-mouth and referral marketing literature [Godes and Mayzlin, 2009, Biyalogorsky et al., 2001, Schmitt et al., 2011, Chae et al., 2016, Van Den Bulte et al., 2018, Kumar et al., 2010]. The two-level experimental design allows causal measurement of the peer effects, an externality that has not received much attention in prior literature and yet can have important managerial implications [Hartmann, 2010]. Extant marketing literature concerning peer effects are mostly in the context of CRM campaigns, which do not have an explicit social component. Those papers use peer encouragement design or cluster randomized trials (e.g. Ascarza et al. 2017 and Godinho de Matos et al. 2016). These approaches require detailed customer network structure for proper randomization. Such network data, however, are not always available to a firm and can be costly to acquire. My approach, in contrast, relaxes the Stable Unit Treatment Value Assumption (SUTVA) without using such network information.

1.4 Institutional Context

This section provides more information on the social media platform WeChat and the online grocery partner.

WeChat is a leading social platform both in China and worldwide.[9] Even leading firms in Silicon Valley closely monitor its product innovations to emulate its successes. A few factors make WeChat an important setting for studying social promotions. First, unlike western social

[9]WeChat's parent company, Tencent Holdings Ltd., overtook Wells Fargo Inc. on April 5th, 2017 to become the world's tenth most valuable company by market value.

media platforms, one can browse any website and directly make purchases within WeChat. Second, unlike the dominant sestern social media platforms, the WeChat news feed favors friends' posts more than sponsored ads. It displays all posts chronologically with no algorithmic filtering or ranking. This gives shared content from friends more visibility.

The e-commerce partner for this study is a stand-alone mid/high-end online grocery shopping site that carries its own inventory. An emerging "unicorn" valued at 785M USD according to Crunchbase, it became profitable in November 2017. It currently delivers to a wide range of cities in China throughout north, east, and south China. Like other fast-growing Chinese tech companies, it is also primarily mobile and around 85% of its GMV comes via mobile.

The experimental deal on the promoted products was only valid during the day of the experiment on Aug 10th, 2017. The coupon's face value is 109 RMB and the original price is 188 RMB ($\approx$ 16 off 27 USD, or 60% off) and the customer could buy the promoted bundle at the price of 79 RMB ($\approx$ 11 USD).[10] The experimental data sample consists of all customers exposed to control or treatment, i.e. (logged-in) customers who opened the experimental deal. From now on, I refer to these customers as the ones who are considering the experimental deal.

1.5 Randomization Checks, Definitions and Notations, and Design Benefits

Randomization Checks

Following the convention in the literature, I verify that randomization is proper both at the city and individual levels [Bruhn and McKenzie, 2009]. I construct a set of pre-treatment covariates and for each one I perform t-test to check if the means are balanced at the 5% significance level. If not, I also test for any potential significant difference in the medians and empirical distributions using the Mann–Whitney U-test and KS-test, resp. I do detect some significant difference on certain covariates under all three tests. In this case, I also run the complementary test for joint orthogonality to ensure that this is expected by chance. More precisely, I test the joint hypothesis that $b_1 = b_2 = \ldots = b_I = 0$ running a linear regression with an F-test and a probit regression with a chi-squared test on the following,

$$Z_i = a + \sum_i b_i X_i + u,$$

where Z_i is the treatment/control indicator and $\{X_i\}$ is the set of pre-treatment covariates.

I present balance tests on city characteristics in Table 1.1 (see Appendix A.5 Table A.8 for the complete list). These include broadly two types, fundamental and firm-specific characteristics. The fundamental characteristics include information on population, GRP, and average income [Chen, 2015]. I construct the firm-specific tests by aggregating individual customer behavior to the city level. I confirm that city-level randomization is proper. For the one out of 37 city characteristics that do not pass the t-test, I verify that medians and empirical distributions are not different at the 10% significance level. This is expected by chance because the joint orthogonality test has a p-value of 0.87 and 0.71 under the F-test and chi-squared test, respectively.

[10] At most one coupon can be applied per order, but a customer may choose to make multiple orders and redeem the same coupon multiple times.

Next, I present balance tests on individual customer characteristics (see Table 1.2). Broadly, I consider three categories: customers' past purchase behavior, experiment-specific measures, and customers' reactions to past marketing activities. First, behavioral characteristics of customers' past purchases include canonical RFM measures in the CPG industry, i.e. Recency, Frequency, and Monetary Value (RFM). Secondly, I examine balance on other pre-treatment experiment-specific measures. These include whether the customer is a new customer who registered during the experiment, whether the customer arrives via referral, and her days since registration. Finally, I look at how customers respond to past marketing activities, including their coupon, discount, and gift card usage. I conclude that individual-level randomization is also proper: all covariates pass the t-test at the 5% significance level.

Table 1.1: City-level balance test

	All	Treatment	Control	t.p	u.p	ks.p
Fundamental City Characteristics						
Household Registered Population of City (10,000 people)	573.702	601.356	545.184	0.483	0.493	0.663
Total Land Area of Administrative Region of City (sq.km)	9,329.985	10,108.640	8,527	0.301	0.607	0.324
GRP of City (10,000 yuan)	51,183,449	51,869,604	50,475,852	0.915	0.974	0.948
Past Purchase Behavior						
Days Since Recent Purchase	167.227	170.822	163.521	0.235	0.588	0.884
Number of Orders Last Year	69,925.540	66,310.640	73,653.410	0.886	0.553	0.654
Percentage of Users Purchased Promoted Products Before	0.064	0.065	0.063	0.681	0.407	0.111
Average Monetary Value per Order	148.300	150.725	145.799	0.678	0.984	0.823
Response to Past Marketing Activities						
Number of Orders Using Coupons Last Year	34,761.390	34,327.640	35,208.690	0.971	0.591	0.880
Number of Orders Using Gift Cards Last Year	17,543.480	16,598.850	18,517.620	0.867	0.718	0.980
Number of Orders with Direct Discounts Last Year	34,555.540	31,799.330	37,397.880	0.832	0.553	0.848
Average Monetary Value of Coupons per Order	31.689	31.783	31.592	0.903	0.477	0.449
Average Monetary Value of Gift Cards per Order	57.703	57.876	57.524	0.968	0.984	0.829
Average Monetary Value of Direct Discounts per Order	11.170	11.477	10.853	0.448	0.799	0.985
Experiment Specific						
Number of Organic Customers	158.846	174.030	143.188	0.773	0.354	0.695
Number of observations	65	33	32			

Notes: Columns (1), (2), and (3) report the means of city characteristics in both treatment and control, the treatment, and control subsamples, respectively. These characteristics include broadly two-types, fundamental characteristics (such as city population, area, and GRP) and others constructed using firm-specific information. Columns (4), (5), and (6) report the p-values of t-test, Mann-Whitney U-test, and Kolmogorov-Smirnov test on whether treated and control cities are significantly different in terms of the mean, median, and empirical distributions of each covariate, respectively. For the balance test results on the complete list of fundamental city characteristics, see Appendix A.5 Table A.8.

Definitions and Notation

For ease of exposition, I summarize the terminologies and notations in Tables 1.3 and 1.4 before further discussion. The customers in my data consist of those who log in and are exposed to the randomized coupon page during the experiment. From now on, unless otherwise noted, this will be my default experimental data sample. Given that these customers already chose to open the coupon page, they have exhibited interest in the promoted products. Therefore, I assume that they have the promoted products in their consideration sets.

Design Benefits

The two-level randomization provides insights into how the promotional intensity affects customer selection. Specifically, I test whether requiring a lower fraction of peers to share results

Table 1.2: Covariates balance test between S and NS customers in 50% NS cities

	Both	NS Customers	S Customers	t.p	u.p	ks.p
Past Purchase Behavior						
Days Since Recent Purchase	59.182	58.599	59.782	0.647	0.357	0.664
Number of Orders Last Year	18.677	19.566	17.772	0.094	0.045	0.006
Purchased Exact Promoted Bundle Before	0.332	0.336	0.328	0.538	0.538	1.000
Purchased Promoted Products Before	0.381	0.385	0.378	0.561	0.561	1.000
Average Monetary Value Per Order	87.167	86.838	87.501	0.739	0.773	0.725
Experiment Specific						
New Customer	0.098	0.093	0.103	0.219	0.219	0.999
Days Since Registration	415.798	412.754	418.898	0.485	0.516	0.250
Referred by Others	0.036	0.034	0.039	0.270	0.270	1
Response to Past Marketing Activities						
Number of Orders Using Coupons Last Year	8.789	9.300	8.269	0.048	0.061	0.132
Number of Orders Using Gift Cards Last Year	2.244	2.502	1.980	0.340	0.523	1
Number of Orders with Direct Discount Last Year	8.830	9.439	8.210	0.066	0.060	0.102
Average Monetary Value of Coupons Per Order	34.629	34.729	34.529	0.841	0.627	0.665
Average Monetary Value of Gift Cards Per Order	10.939	10.181	11.711	0.201	0.439	0.848
Average Monetary Value of Direct Discount Per Order	9.331	9.363	9.299	0.888	0.243	0.284
Average Monetary Value of Coupons and Direct Discount Per Order	43.960	44.091	43.827	0.812	0.639	0.859
City Characteristics						
First-Tier Cities	0.664	0.669	0.660	0.479	0.479	1.000
Number of observations	5,959	3,006	2,953			

Notes: Columns (1), (2), and (3) report the means of individual characteristics in both treatment and control, the treatment, and control subsamples, respectively. Columns (4), (5), and (6) report the p-values of t-test, Mann-Whitney U-test, and Kolmogorov-Smirnov test on whether S and NS customers are significantly different in terms of the mean, median, and empirical distributions of each covariate, respectively.

in a different set of customers in considering the promoted products (i.e. whether they open the experimental coupon page after logging in). I find that compared with mandating a a random half of the customers to share, mandating all to share induces the following types of customers to consider the promoted products: those who have a larger average order size, who have purchased the promoted products before, who used more gift cards, and who have been registered for a longer period of time (see Appendix Table A.1 and Appendix Figure A.4).[11]

1.6 Average Treatment Effects - Intensive Margins on the Promoted Products and Extensive Margins

Before presenting the model used for measuring the overall social impact, I first present results on the average treatment effects, which constitute key parameters in that model.

To identify the direct treatment effects of interest, the intensive and the extensive margins, I estimate regressions of the following forms on customers within treated cities:

$$Y_i^I = r_{\mathrm{S}}(50\%) + i(\alpha) Z_i + X_i^\top w^I + \varepsilon_i^I,$$

and

$$Y_i^E = s(\alpha = 50\%) - \mathrm{Ext}^\alpha Z_i + X_i^\top w^E + \varepsilon_i^E.$$

The intensive margin outcomes (Y_i^I's) include customer i's purchase decision, expenditure, and profits for the promoted products (i.e. Revenue - Costs of Goods Sold (COGS)). The

[11]While it would be ideal to shed more light on the selection process with additional Chinese customer demographics information, such data are hard to obtain. A thorough search of library databases of both American and Chinese universities and correspondences with other Chinese scholars and industry practitioners, such as AC Nielson CEO, suggest that the most granular Chinese consumer demographics data is at the city-level.

Table 1.3: List of terminologies

Terminologies	Definitions
S customer	A control customer who needs to share to obtain a deal.
NS customer	A treated customer who can redeem the deal directly.
referral/lead/child customer	A customer who opens the deal shared by a friend on WeChat.
parent customer	A customer who introduces at least one child customer.
(promotional) intensity	The fraction of S customers in a city.

Table 1.4: List of notations

Notations	Terminologies	Definitions
Ω	The set of organic customers	Customers who arrive directly at the website, i.e. non-children customers.
$\#\Omega$	The number of organic customers	
γ	The type of product	If omitted or $\gamma = p$ stands for promoted products; $\gamma = a$ stands for any product.
R^γ	The total revenue of product γ	
α	The treatment intensity	The fraction of NS customers in a city. An α NS city is one where $1-\alpha$ of customers are randomly assigned to S. Each control city is a 0% NS city and each treated city is a 50% NS city.
$s(\alpha)$	Social multiplier	The ratio of referral customers to S customers in an α NS city, $\mathrm{s}(\alpha) = \frac{\#\text{ referral customers}}{(1-\alpha)\#\Omega}$.
$r^\gamma_{\mathrm{S}}(\alpha)$, $r^\gamma_{\mathrm{NS}}(\alpha)$	Average revenue	The average revenue of product γ among S or NS customers from an α NS city.
$i(\alpha)$	Intensive margin	Holding an individual customer fixed, the change in her intensive margin outcome, such as revenue either from the exact promoted bundle or from any product, due to removing the sharing requirement. $\mathrm{i}^\gamma (= i^\gamma(\alpha)) = r^\gamma_{\mathrm{NS}}(\alpha) - r^\gamma_{\mathrm{S}}(\alpha)$
IntSpl	Peer effects on the intensive margin	Holding an S customer fixed, the change in her spending on a product when her city is an α_1 NS city versus an α_2 NS city. $\mathrm{IntSpl} = r^\gamma_{\mathrm{S}}(\alpha_1) - r^\gamma_{\mathrm{S}}(\alpha_2)$
X_i	Customer characteristics	Customer i's observed characteristics.
Z_i	Treatment Status	$Z_i = 0,1$ denotes customer i's assignment to S and NS, respectively.
P_i	Profit from customer i herself	Under the potential outcomes framework, i's observed profit is $P_i = P_i(Z_i = z_i)$.
$PChild_i$	Profit from i's children	The profit customer i contributes indirectly from her children.

extensive margin outcomes (Y_i^E's) include the number of children customers a focal customer produces, as well as the number of new, existing, and buying customers among the children customers. Throughout the paper, I use the number of children customers introduced as the primary representative outcome for the extensive margin of interest. Even if a child customer comes and does not make any immediate purchase, the social coupon may still contribute toward increasing her future purchase intent through increased brand awareness.

I use Z_i to denote the treatment indicator, where a treated customer is assigned a regular NS coupon with no social incentive. In some specifications, for potential improvement in the precision of the estimates, I include covariates to control for unobserved heterogeneity.

Intensive Margins on Promoted Products

I now present the direct treatment effects of interest, $i^p(\alpha = 50\%)$. For each outcome, I present specifications estimated using OLS (see Table 1.5). Removing the sharing requirement significantly boosts a customer's propensity to purchase the promoted products and the associated expenditure. I find that, within a treated city (where a random half of the customers considering the experimental deal are forced to share), removing the sharing requirement increases a customer's average propensity to buy the promoted products from 28% to 34%. This 21 percent (6%/28%) increase is economically and statistically significant.On average, removing the sharing requirement increases a customer's expenditure on the promoted products by 0.7 USD (or 4.4 RMB), which corresponds to a 20.6% increase. Firms' profit from these promoted products' sales, however, does not change significantly due to cost variability. The margin for each promoted product is relatively thin, and the cost distribution varies substantially across different cities due to different wholesale costs in different cities (see Appendix A.3 Figure A.7 for the empirical cumulative distribution plot of the promoted products' profitability).

Moreover, these point estimates can also be interpreted as the difference in means when comparing an NS customer in a 50% NS city with an S customer in a 0% NS city. Here, one may raise the selectivity concern that these two groups of customers are quite different. I argue that substantively this is not a concern. First, given that the coefficient of interest does not vary much among the baseline and alternative specifications, not controlling for selection on observables does not lead to omitted variable bias of substantive concern. Furthermore, assuming that I observe a random subset of all customer characteristics, I verify that selection on unobservables is also highly unlikely to alter the results. The idea here is that selection on observables provides guidance for the amount of selection on unobservables; I refer interested readers to Appendix A.1.

Extensive Margins - Social Promotions as Effective New Customer Acquisition and Existing Customer Retention Strategies

This section presents findings on the extensive margins. More precisely, I investigate how many referrals a customer generates when she is versus is not forced to share within the "50% sharing" cities.

I find statistical and economic significance in the extensive margins. Within a 50% NS city, each arriving customer assigned to the S condition generates, on average, 0.08 units of children customers who also consider buying the promoted products (95% confidence interval of .05 to .11 units; see Table 1.5). In dollar values, this extensive margin corresponds to an average profit contribution of 0.033 USD (95% confidence interval of 0.020 to 0.045 USD or 0.14 to 0.32 RMB). In the long run, this corresponds with a projected average lifetime value of 117 USD (or 819 RMB) per customer. As one would expect, NS customers rarely share and produce almost no children customers.

Moreover, social promotions contribute toward both new customer acquisition and existing customer retention. When forced to share, a customer brings on average an additional 0.03 and 0.04 units of new and existing customers, respectively. In dollar values, these correspond to average profit contributions of 0.016 USD (or 0.112 RMB) and 0.021 USD (or 0.149 RMB), and projected average lifetime values of 62 USD (or 434 RMB) and 83 USD (or 579 RMB), respectively.

Table 1.5: Average treatment effects on intensive and extensive margins

	Across Customer Effects			Within Customer Effects on Promoted Products	
	# Customers Introduced	# New Customers Introduced	# Existing Customers Introduced	Promoted Products Purchase Propensity	Promoted Products Expenditure
	(1)	(2)	(3)	(4)	(5)
Assigned to Social	0.081^{***}	0.039^{***}	0.042^{***}	-0.057^{***}	-4.530^{***}
	(0.014)	(0.006)	(0.008)	(0.010)	(0.865)
Constant	0.001	0.001	0.000	0.337^{***}	26.296^{***}
	(0.001)	(0.001)	(0.000)	(0.013)	(1.348)
Observations	5,959	5,959	5,959	5,959	5,959
R^2	0.017	0.010	0.012	0.004	0.003
adjusted R^2	0.017	0.010	0.011	0.004	0.003

Notes: These regressions are restricted to 50% NS cities only. The dependent variables in column (1) - (5) are number of customers introduced, number of new customers introduced, number of existing customers introduced, promoted product purchase propensity, and promoted product expenditure respectively. Standard errors are robust to heteroskedasticity and clustered at the city level.

1.7 The Delta Method - Evaluating the Overall Social Impact on Promoted Products

Now that I have addressed the average treatment effects on the intensive and extensive margins respectively, I discuss how to assess the trade-off between these two margins. To demonstrate how to evaluate overall efficacy, I compare the following two social promotional strategies: forcing all versus a random half of the customers to share. To capture the trade-off between the costs and benefits, it is unfeasible to apply a standard regression model on an individual-level outcome directly. On the one hand, the closest metric that comes to mind is an effective final outcome that reflects both a customer's own purchase decision, and, were she a parent, her children's purchase decisions. However, applying a regression model using this effective outcome on all experimental observations will result in double counting those children customers' purchases. On the other hand, leaving out the subset of children customers makes the input data unbalanced in the first place. Therefore, to answer the aggregate-level question, it is necessary to develop a model that examines the outcome as a function of the promotional intensity.

This section proceeds as follows. First, I discuss the model set-up. Second, I present results applying the model to evaluate effectiveness on the promoted products. Third, I illustrate how aggregate-level business decisions could be mislead in the wrong direction if the no-peer effects assumption fails to hold.

Model

To simplify the analysis, the model is premised on the following assumptions that are empirically validated in the setting of the field study.

NS customers bring no leads. This implies that $s(50\%) = \text{Ext}^{50\%}$.[12]

There is only one generation of children and the random assignment of children customer does not significantly change α.[13]

No-Peer Effects Assumption The average response of S customers does not depend on whether all or a random half of the peers in the same market are required to share. There are no peer effects on both the intensive and extensive margins, i.e.

$$r_S^\gamma(\alpha) \equiv r_S^\gamma \qquad \text{and} \qquad s(\alpha) \equiv s,^{14}$$

where $r_S^\gamma(\alpha)$ denote the average revenue among S customers, respectively, γ denotes the type of product, and $s(1-\alpha)$ denotes the social multiplier, i.e. the average number of referrals generated by an S customer in an $(1-\alpha)$ NS city.

No-Cannibalization Assumption The promotional intensity α does not affect the number of organic or non-referral customers, i.e.

$$\#\Omega^{\alpha=50\%} = \#\Omega^{\alpha=0\%}$$

Under Assumption 1.7 (empirically validated, see last row of Table 1.1), to evaluate overall effectiveness of the social promotions, it suffices to conduct inference on an individual-level effective outcome. This effective measure properly captures the number of referrals a focal customer generates. I present below a model in which the aggregate outcome of interest is total revenue; the corresponding individual-level outcome of interest is average spending per customer. Similarly, to examine alternative outcomes, such as demand or profitability, it suffices to replace $R^\gamma(\alpha)$ and $r^\gamma(\alpha)$'s with the corresponding metrics of interest.

Given that the portion of revenue from customers outside of the experiment does not significantly differ among 50% NS and 0% NS cities, it suffices to consider total revenue from customers in the experiment. The total revenue of promoted products (or respectively of all products) among customers considering the experimental deal in a city with intensity α is

$$R^\gamma(\alpha) = R(\vec{\beta}^{\alpha,\gamma}) = (\#\Omega + (1-\alpha)\#\Omega \cdot s(\alpha))(\alpha r_{NS}^\gamma(\alpha) + (1-\alpha) r_S^\gamma(\alpha)) \tag{1.1}$$
$$= (\#\Omega + (1-\alpha)\#\Omega \cdot s(\alpha))(\alpha(r_S^\gamma(\alpha) + i^\gamma(\alpha)) + (1-\alpha) r_S^\gamma(\alpha)) \tag{1.2}$$

[12]Even though NS customers may also choose to voluntarily share, this rarely happens as expected. A total of 3006 NS customers only bring 3 leads to the site, which is negligible.

[13]The number of referrals with degrees ≥ 2 in the experiment is negligible, where degree refers to the referral customer's number of generation. If there are many higher-degree referrals, one can also attribute them to the organic customers correspondingly. Additionally, note that even if an existing customer is referred to the deal page again, he or she will be reassigned to either group with the corresponding treatment intensity of that city

[14]This assumption is experiment specific. In our experiment, S customers in 50% NS cities and S customers in 0% NS cities have the same average propensity to purchase promoted products

where $\vec{\beta}^{\alpha,\gamma} = \begin{pmatrix} s(\alpha) \\ r^{\gamma}_{\mathrm{S}}(\alpha) \\ i^{\gamma}(\alpha) \end{pmatrix}$, and R^{γ} denotes the total revenue of product γ.

Then, an individual customer's effective spending on the promoted products (or any products) conditional on landing the experimental deal page is

$$r^{\gamma}_{e}(\alpha) = \frac{R^{\gamma}(\alpha)}{\#\Omega} = (1 + (1-\alpha)\cdot s(\alpha))(\alpha r^{\gamma}_{\mathrm{NS}}(\alpha) + (1-\alpha) r^{\gamma}_{\mathrm{S}}(\alpha)),$$

where I explicitly take into account additional purchases from children customers.

Then, to assess the overall campaign effectiveness on sales of the promoted products, I test the alternative hypothesis

$$H_A : r_e(\alpha_1 = 50\%) - r_e(\alpha_2 = 0\%) \neq 0,$$

against the null hypothesis

$$H_0 : r_e(\alpha_1 = 50\%) - r_e(\alpha_2 = 0\%) = 0.$$

Let

$$\Delta r_e(\alpha_1, \alpha_2) = r_e(\alpha_1) - r_e(\alpha_2).$$

I construct a 95% confidence interval for

$$\Delta r_e(\alpha_1 = 50\%, \alpha_2 = 0\%) = r_e(\alpha_1 = 50\%) - r_e(\alpha_2 = 0\%).$$ [15]

The Overall Social Impact on Promoted Products

I find that for a given customer in the experiment, when the firm randomly removes the sharing requirement with a 50% probability, her effective total expenditure on these promoted products increases significantly by 5% (0.2 USD or 1.4 RMB and a 95% confidence interval of [0.1,0.3] USD or [0.6, 2.4] RMB).

Discussion on Managerial Relevance of the Peer Effects

This subsection illustrates how failing to account for sufficiently large peer effects may result in a sub-optimal aggregate business decision. In this case, the decision of interest is to choose the fraction of customers on whom to impose the sharing requirement. Assume that there is no peer effect on the extensive margin, I provide an example on how to compute the relevant threshold for the intensive margin peer effect. To be precise, I replace Assumption 1.7 with the following.

There could be a peer effect on the intensive margin

$$\mathrm{IntSpl} = r_{\mathrm{S}}(0\%) - r_{\mathrm{S}}(50\%).$$

[15] Specifically, I assume that

$$\sqrt{n}(\widehat{\vec{\beta}_{\alpha,\gamma}} - \vec{\beta}_{\alpha,\gamma}) \xrightarrow{d} \xi \sim N(0, V).$$

and apply the delta method. The computation details are included in Appendix A.2.

There is no peer effect on the extensive margin, i.e., for any $\alpha \in \{0\%, 50\%\}, s(\alpha) \equiv s$. Subsection 1.7 concludes that overall, removing the sharing requirement with a 50% probability at random results in higher sales of the promoted products. This subsection shows that the conclusion will reverse if the peer effect is sufficiently large.
Recall

$$r_e(0\%) = (1+s)r_S(0\%) = (1+s)(r_S(50\%) + \text{IntSpl}),$$

$$r_e(50\%) = (1+0.5\cdot s)(r_S(50\%) + 0.5\cdot i(50\%)),$$

and obtaining the opposite aggregate-level result means

$$r_e(0\%) - r_e(50\%) > 0.$$

Therefore, substituting in s, $r_S(50\%)$, and $i(50\%)$ gives a threshold of IntSpl > 0.197 USD. That is, if the peer effect on average revenue of promoted products is greater than 0.197 USD, the conclusion in last subsection will reverse (see Table 1.6 for thresholds on other intensive margin outcomes). For example, if the peer effect on the propensity to buy the promoted products is greater than 0.017, this will also lead to a sub-optimal managerial decision.

To make sense of the peer effect, I normalize it by the direct effect of the intensive margin and look at its relative size. Since removing the social incentive increases a customer's purchase probability by 5.7%, then the managerially relevant threshold for the peer effect is approximately 29.7% (0.017/0.057).

Table 1.6: Thresholds for peer effects of managerial relevance

Dependent Variable	Threshold for Peer Effects (Average Among S customers in 50% NS Cities - Average Among S customers in 0% NS Cities)
propensity to buy Promoted Products	0.017
Expenditure on Promoted Products	\$0.197
Profit from Promoted Products	\$0.002

Notes: This table reports the thresholds for peer effects on intensive margin outcomes. If the magnitude exceeds the threshold, failing to account for the peer effects can result in sub-optimal choice of social promotional strategy.

1.8 The Overall Social Impact on All Products - Consumer Spending and Total Demand

The results on promoted products beg the question of if the pain associated with sharing outweighs the word-of-mouth benefits. Managerial decision making requires more than the intermediate conclusion that only concerns the promoted products. The social promotions may also have spillover effects on sales of the non-promoted products. This section presents the results on the overall social impact regarding all products on the platform.

Overall, requiring all customers to share is significantly better than requiring a random half of the customers to share. A customer's propensity to buy any products from the platform, the corresponding expenditure, and profit generated do not significantly differ when required versus not required to share. When forced to share, some customers substitute and purchase other products from the same grocery site; these are mostly fruits.

Substitution Patterns

The firm carries grocery products in multiple categories, including fruits, snacks, maternal-infant care products, alcohol, seafood, and meat. A natural question is how the social incentive affects sales of the non-promoted products.

Figure 1.1 presents a summary diagram on how the social incentive affects customers' propensity to buy different kinds of products. First, the customer basket composition does not significantly differ for S customers in 0% NS cities vs S customers in 50% NS cities. A comparison of the first two columns in Figure 1.1 illustrates this point. As a reminder, I have also established in my earlier presentation that the probability of buying the promoted products alone is similar. When the sharing requirement is removed, however, more consumers who would have only bought other non-promoted products instead buy the promoted products (a wider distance between the red and blue dots in column 3 than that in column 2). The distance between the red and blue dots represents the fraction of customers who buy the promoted products. They may or may not buy other products together with the promoted products. Secondly, the sharing requirement mostly induces substitution within the fruits category. The pattern remains robust when I examine the outcome expenditure instead (see Appendix Table A.9). The fraction of customers who only buy other non-fruit products remains stable across S and NS treatment conditions (the parallel blue and green lines in the last two columns). That is, if a customer only consumes non-fruits, such as wine, she would still only buy non-fruits, regardless of whether she sees an S or NS promoted products coupon. The most popular substitutes for the promoted products include other similar exotic imported fruits: Durian Monthons, American pluots, and New Zealand's kiwifruits; their sales constitute 29.27%, 19.98%, and 6.45%, resp., of fruits consumption on non-promoted products.

Figure 1.1: Substitution patterns

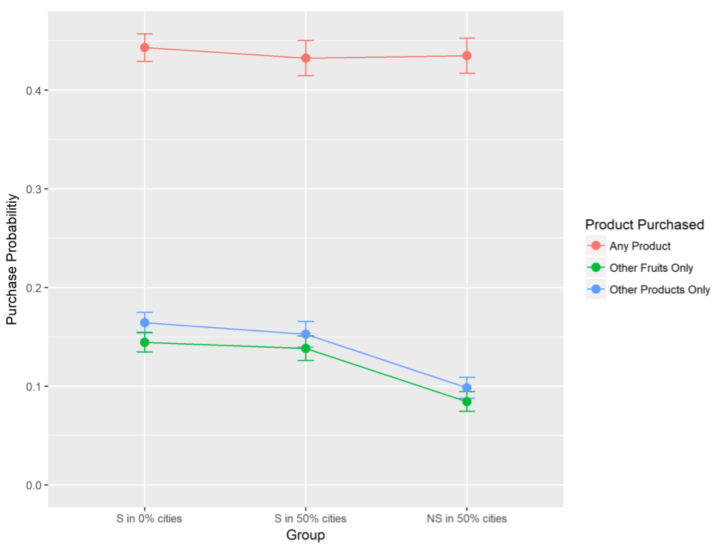

Notes: This figure shows a few average purchase probabilities and the corresponding 2 standard deviations confidence bands to illustrate consumer substitution patterns. Red stands for the propensity to buy any product, blue for any non-promoted products, and green for any fruit other than the promoted products. The distance between the red and blue dots in each column represents the average probability of buying the promoted products; the distance between the blue and green dots in each column represents the average probability of buying exclusively non-fruits.

The Overall Social Impact on All Products

Because the average treatment effects on the intensive and extensive margins suggest that high social intensity (0 % NS) is better than low social intensity (50 % NS), I test the following alternative hypothesis

$$H_A : r_e^a(\alpha = 50\%) - r_e^a(\alpha = 0\%) \neq 0.$$

against the null hypothesis

$$H_0 : r_e^a(\alpha = 50\%) - r_e^a(\alpha = 0\%) = 0.$$

The point estimate for $\Delta r_e^a(\alpha_1 = 50\%, \alpha_2 = 0\%)$ is -0.3 USD or -2 RMB, and it is statistically significant with a p-value of 0.036 and a 95% confidence interval of $[-.5, -.1]$ USD. Namely, conditional on a customer landing on the promoted products' coupon page, if she belongs to a 50% NS city, her effective spending on any products (accounting for the children sales she produces) decreases significantly by 0.3 USD from 6.8 USD to 6.5 USD on average, compared with if she were in a 0% NS city (see Figure 1.1 for the baseline conversion rates before adjusting for children sales). This corresponds to a 4 percent decrease (a 95% confidence interval of 1 to 7 percent). When I examine the purchase propensity instead, I find that removing the sharing requirement randomly for half of the customers results in a statistically and economically significant drop. The loss is about 2% for each arriving customer (a 95% confidence interval of [-0.03 to -0.01]). Table 1.7 summarizes the point estimates of $\Delta r_e^\gamma(\alpha_1 = 50\%, \alpha_2 = 0\%)$ for different outcomes.

Table 1.7: Delta method summary

Dependent Variable	$r_e^\gamma(\alpha_1 = 50\%)$	$r_e^\gamma(\alpha_2 = 0\%)$	$\Delta r_e^\gamma(\alpha_1 = 50\%, \alpha_2 = 0\%)$	p-value
Propensity to buy Promoted Products	0.3201	0.3004	0.0197	.001***
Expenditure on Promoted Products	\$3.5676	\$3.3481	\$0.2196	.001***
Propensity to buy All Products	0.4440	0.4640	−0.0199	.001***
Expenditure on All Products	\$6.4826	\$6.7743	−\$0.2917	.036*

*p<0.1; **p<0.05; ***p<0.01

Notes: This table summarizes the estimates of $r_e^\gamma(\alpha_1 = 50\%)$, $r_e^\gamma(\alpha_2 = 0\%)$, $\Delta r_e^\gamma(\alpha_1 = 50\%, \alpha_2 = 0\%)$ for different intensive margin outcomes.

1.9 Customer Heterogeneity - Propensity to Share and Extensive Margins

So far the analysis on optimal promotional intensity only relies on identification of the average treatment effects. A natural next step is to leverage additional information to celebrate and analyze customer heterogeneity to inform the analysis.

Empirical Strategy

I first present my empirical strategy for analyzing observed customer heterogeneity. Following the literature convention, I use the linear model with interaction term as the primary specification. I divide the customers into segments based on each covariate of interest and test if the conditional average treatment effects are significantly different across segments. More precisely, I divide customers into different categories if the covariate of interest is a categorical variable or divide by the median if the covariate is continuous (by-median model).[16] Generally, to verify robustness of the results, I test two alternative models. First, I test for heterogeneity on the linear model specification where I segment the customers by mean of the covariate of interest (by-mean model). Second, to further relax the sparsity assumption and deal with potential issues of small cell sizes, I apply a recently developed machine learning-based model, an instance of generalized random forests (grf), as yet another alternative specification to test for robustness of the results [Athey et al., 2016]. Grf can be particularly helpful if a firm wishes to construct additional customer characteristics capturing consumer search, browsing, or product review behavior that I do not explicitly model here.

The heterogeneity results I present are statistically significant in the primary specification and also satisfy at least one of the following robustness criteria: a) the results remain consistent and statistically significant in the by-mean model; b) among the 16 input covariates, the covariate of interest is ranked among the 10 most important ones based on grf's variable importance measure.[17]

For each heterogeneity result under the primary specification, I present the conditional average treatment effects (CATEs) estimates on both segments (see Table 1.8). The baseline includes no additional controls other than the covariate I test heterogeneity on. The full specification (in appendix) includes a set of covariates that could potentially lead to more precise estimates. For example, when I test heterogeneity on recency, the corresponding interaction-terms specifications are as follows,

$$Y_i = \beta_0 + \beta_1 Recency_i + \beta_2 Z_i + \beta_3 Recency_i * Z_i + \varepsilon_i$$
$$Y_i = \beta_0 + \beta_1 Recency_i + \beta_2 Z_i + \beta_3 Recency_i * Z_i + X_i^T \gamma + \varepsilon_i,$$

where $Recency_i$ is an indicator of if the customer's days since recent purchase is larger than the median. Following the literature convention, I include additional covariates X_i^T in the full specification that may lead to potentially more precise estimates of the CATEs. Specifically, I run lasso on the average treatment effects specification and include the selected set of covariates.

Heterogeneous treatment effects estimation via machine learning-based (ML-based) methods has received considerable attention lately. Grf is one such method from this family of ML-based methods that have started to gain traction among applied economists. Without imposing any sparsity assumption on the true data generating process, grf provides a post-hoc data-driven approach to mine for heterogeneity in causal effects with valid standard errors. It nests my primary specification, the conventional interaction-terms specification, as a special case. It alleviates the potential over-fitting concern and more flexibly captures non-linear structure in the

[16] One might raise the concern of overrejection of the null hypothesis, given that I am testing many dimensions of heterogeneity. First, managerially speaking, false discovery is not so much a concern. Firm would be better off having some false positives than false negatives as additional experimentation based on the false positives will always inject new input into the feedback loop. Secondly, the grf robustness check itself serves as a data-driven search of heterogeneity, which also alleviates this multiple-comparison concern.

[17] I also verify that for covariates that are only significant under by-mean model or only ranked among the 10 most important ones in grf but not significant in the by-median model, these heterogeneity findings are substantively consistent with what shows significance under the by-median model.

data. Compared with the commonly-used interaction-terms specification, grf is robust to mis-specification and alleviates concerns of ex-post data-mining by allowing the econometrician to specify a set of potential covariates.

Another advantage of grf is that it quantifies relative importance of each covariate. The variable importance measure, roughly speaking, reflects how much information gain one obtains from adding a covariate to the model versus not adding it. More precisely, it reflects a weighted average of how frequently the covariate is used for node splitting when constructing each individual causal tree that forms the forest ensemble.

To the best of my knowledge, grf is also a more conservative way of testing for heterogeneity; several recent quantitative marketing applications find that grf could be over-smoothing and therefore might overlook some heterogeneity [Guo et al., 2017, Ebrahim et al., Shah et al., 2019]. Appendix A.3 Figure A.8 presents one such example that compares the decile plot of CATE estimated by interaction-terms model and grf for the frequency covariate. I refer the interested reader to Athey et al. [2016] for the technical details.

For brevity, for each category of heterogeneity, I present regression estimates from the primary by-median specification. Specifically, I present results on the following two outcomes, an indicator of if the customer successfully shares the coupon, and a representative outcome for the extensive margin, the number of customers introduced. Appendix A.4. presents similar results using other outcomes, (i.e. the number of existing customers introduced, the number of new customers introduced, and the number of buyers introduced), and robustness checks using the more flexible non-linear grf and by-mean specifications. When a certain type of customers generate more referrals, they most likely also have significantly higher propensity to share. In particular, the primary outcome within the extensive margin family, i.e. the number of customers introduced, is fairly representative for other extensive margin outcomes; more productive parents generally contribute significantly more toward existing customer retention, new customer acquisition, and additional buyer production. As examples, I discuss heterogeneity results on the full spectrum of outcomes in the city-level heterogeneity subsection and days since registration subsection. For brevity, I only present estimates on the primary outcome for remaining discussion on other categories of heterogeneity. Generally, at least each covariate-outcome pair shows statistical significance under the linear by-median or by-mean model or shows importance in grf.

Construction of Observed Customer Characteristics

Broadly, I test heterogeneity on observed variables capturing customer value, price sensitivity, and other ad-hoc customer characteristics, such as if a customer is from a first-tier city, her knowledge of the promoted products quality, and gift card usage (see Table 1.2 for the list of covariates; the only one left out is the Principal Component Score for the city the customer is from[18]).

[18]Given that the number of city characteristics is relatively large, I use factor analysis for dimensionality reduction. The Principal Component Score refers to the score of the first principal component (PC_1), which accounts for 98.2% of the variation.

Summary of Findings on Heterogeneity

I find heterogeneity at both the city and individual level for a customer's propensity to share and lead-generation potential. At the city level, while requiring sharing helps with new customer acquisition and existing customer retention in both first-tier and non-first-tier cities, the contribution are significantly higher in first-tier than non-first-tier cities. This may suggest that consumer purchasing power and the breadth of the network may matter more than the strengths of the social ties. At the individual level, constructing characteristics based on customers' preferences and behavior, I find that the following types of customers tend to have higher propensity to share and lead-generation potential: those who are more price sensitive, who have been registered for a longer period of time, who have consumed the promoted products before, who have used gift cards less in the past, who have made a recent purchase longer ago, and who have purchased less frequently in the past year. (see Figures 1.2 and 1.3 for model-free evidence).

Figure 1.2: Model-free evidence on the outcome number of customers introduced

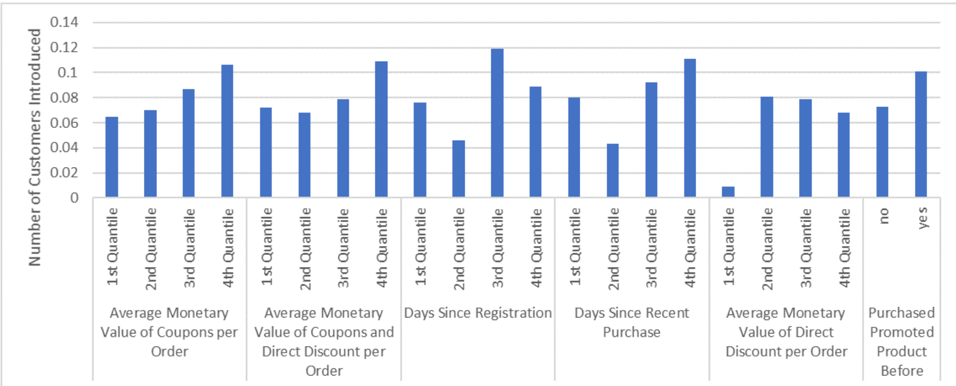

Figure 1.3: Model-free evidence on the outcome number of customers introduced

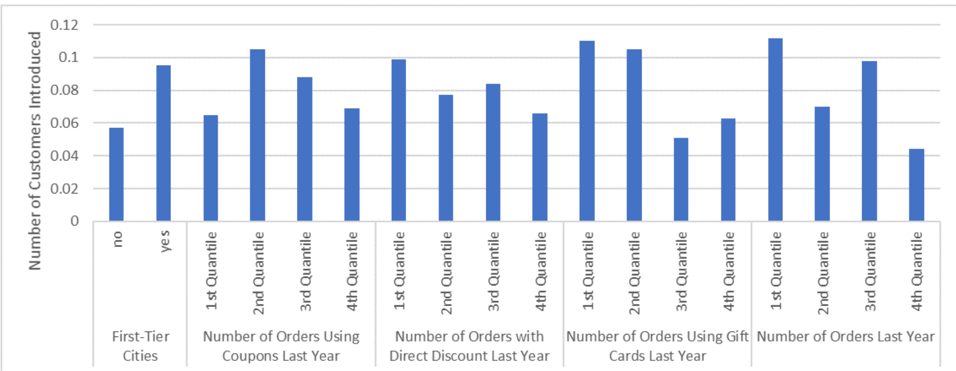

City-level Heterogeneity

One question that piques managerial interest is that, as a growth strategy, for what sort of markets are social promotions better suited and more powerful. For example, the miracle social e-commerce start-up PDD establishes its viral growth through exploiting social connections among customers in small cities (primarily Tier 3 and Tier 4 cities). Unsurprisingly, in my case, a mid/high-end grocery e-commerce company's social promotions turn out more effective among customers from first-tier rather than non-first-tier cities. Ex-ante, it is unclear if social promotions will be more or less effective among customers from smaller versus bigger cities. On the one hand, they might work better among big-city customers for a variety of reasons. First, these people tend to have broader networks, which will result in more visibility for the shared content. Secondly, these people on average demand higher life quality, have higher income and stronger purchasing power. Therefore, big cities enjoy a larger potential customer base that fits the firm's brand positioning and emphasis on high-quality and healthy grocery. Furthermore, ceteris paribus, they might naturally have a higher probability of adopting a relatively expensive and fashionable fruit like like the promoted products. However, on the other hand, it could also be the case that social promotions are more effective among customers from smaller cities. Despite a narrower network, these customers may have potentially stronger social ties. Therefore, conditional on seeing a shared coupon, friends may be more likely to convert. I provide suggestive evidence on how consumers from first-tier cities differ from those from non-first-tier cities.

Specifically, compared with customers from non-first-tier cities, I test if those from first-tier cities a) are more sensitive to the sharing requirement and b) have stronger potential for lead generation. The heterogeneity is statistically and economically significant. I find that customers from first-tier cities tend to generate more leads and have higher propensity to share (see Table 1.8). Each arriving customer in a first-tier city generates an additional 0.09 unit of leads. This is 50% more compared with those from non-first-tier cities, where each arriving customer only generates an additional 0.06 unit of leads. Furthermore, these CATEs are highly statistically significant for customers from both first-tier and non-first-tier cities. Moreover, I also find that social promotions are effective new customer acquisition and existing customer retention strategies in both first-tier and non-first-tier cities. It remains the case that customers from first-tier cities bring in both more new customers and more existing customers; these results are marginally significantly. The point estimates for both new and existing customers are about 0.05 for first-tier, and 0.03 units for non-first-tier cities (with p-values of 0.07 and 0.09, resp., see Table 1.8). One limitation of this result is its external validity; social promotional effectiveness depends on a firm's brand positioning as well as the specific markets it operates in.

Price Sensitivity

Another interesting dimension of heterogeneity is how lead-generation potential differs by customers' price sensitivity. I find that more price-sensitive customers are more productive parents and have higher propensity to share. This is reasonable because these price sensitive customers may have relatively low sharing costs and many price-sensitive friends who are interested in following suit to obtain the discounts. As a matter of fact, there are many WeChat groups formed just for the sake of information sharing on promotions.

Broadly speaking, the price sensitivity measures I test include a customer's responses to past marketing activities based on her past coupon and direct discounts usage, as well as the average order size (i.e. the canonical monetary value measure in retail marketing). More precisely, I capture a customer's degree of price sensitivity using the following covariates: the number of orders using coupons last year, the number of orders using direct discounts last year, the average monetary value of coupons per order, the average monetary value of direct discounts per order, and the average monetary value of coupons and direct discounts per order. Coupons and direct discounts are the two major forms of promotional vehicles the firm offers for customers. I assume that those who adopt promotions more frequently (operationalized as the number of orders using promotions in the past year) and use larger dollar values of promotions are more price sensitive. Furthermore, I presume that customers with smaller average order size are more likely to have lower purchase power and therefore may be more price sensitive.

The finding supports the hypothesis. Indeed, as predicted, those who apply more dollar value of coupons per order on average tend to have stronger lead-generation potential and higher propensity to share (see Tables 1.8, and Appendix A.4 Tables A.2, A.3, and A.4; for other extensive margin outcomes see Appendix A.4 Tables A.5, A.6, and A.7). Results are similar when I examine the average monetary value of coupons and direct discounts per order instead (see Table 1.8, and Appendix A.4 Tables A.2, A.3, and A.4; for other extensive margin outcomes see Appendix A.4 Tables A.2, A.5, A.6, and A.7). Furthermore, it is worth noticing that, compared with the frequency measures, those average monetary value measures generally better capture the heterogeneity. The frequency measures consistently rank less important in grf (see for example, Appendix A.4 Table A.3).

Additionally, I find suggestive evidence of customer inertia; customers tend to stick with using a specific type of promotion.

Those who used more direct discounts in the past have lower propensity to share. It is much more convenient for a customer to apply direct discounts compared with using a coupon. While a simple click on the item page suffices for direct discounts, one has to go through a redemption process if one wishes to enjoy a deal from a coupon. Unsurprisingly, I find that those customers who have used direct discounts less frequently in the past have significantly higher propensity to share (52% versus 48% with a p-value of .03, see Table 1.8 and Appendix A.4 Table A.2).

Table 1.8: Heterogeneity on number of customers introduced and propensity to share

		Number of Customers Introduced			Propensity to Share		
		CATE Greater than Median if Continuous			CATE Greater than Median if Continuous		
		Yes	No	p-value of test (1) = (2)	Yes	No	p-value of test (4) = (5)
Covariate	Median	(1)	(2)	(3)	(4)	(5)	(6)
First-Tier Cities		-0.093***	-0.057***	0.031	-0.522***	-0.460***	0.002
		(0.011)	(0.009)		(0.011)	(0.016)	
Days Since Registration	390	-0.101***	-0.061***	0.011	-0.540***	-0.462***	0.000
		(0.013)	(0.009)		(0.013)	(0.013)	
Average Monetary Value of Coupons per Order	27	-0.097***	-0.065***	0.047	-0.538***	-0.464***	0.000
		(0.013)	(0.010)		(0.013)	(0.013)	
Average Monetary Value of Coupons and Direct Discount per Order	38	—	—	—	-0.526***	-0.476***	0.007
		—	—		(0.013)	(0.013)	
Average Monetary Value per Order	99	-0.060***	-0.098***	0.031	—	—	—
		(0.010)	(0.014)		—	—	
Number of Orders with Direct Discount Last Year	1	—	—	—	-0.480***	-0.519***	0.034
		—	—		(0.013)	(0.013)	
Purchased Promoted Products Before		-0.101***	-0.071***	0.047	-0.579***	-0.463***	0.000
		(0.015)	(0.009)		(0.016)	(0.011)	
Average Monetary Value of Gift Cards per Order	0	-0.034***	-0.089***	0.018	-0.449***	-0.509***	0.025
		(0.010)	(0.009)		(0.025)	(0.010)	
Days Since Recent Purchase	20	-0.106***	-0.052***	0.002	-0.563***	-0.416***	0.000
		(0.015)	(0.009)		(0.015)	(0.014)	
Number of Orders Last Year	10	-0.061***	-0.095***	0.053	-0.456***	-0.519***	0.002
		(0.009)	(0.014)		(0.014)	(0.015)	
Active Customers		—	—	—	-0.490***	-0.542***	0.022
		—	—		(0.010)	(0.021)	

Notes: This table presents the heterogeneity on number of customers introduced and propensity to share. For each covariate, columns (1) and (2) report conditional average treatment effects (CATEs) estimated from regressions without controls (i.e. comparisons of means for S and NS customers in 50% NS cities) on the outcome number of customers introduced. For example, when the covariate of interest is recency, in columns (1) and (2), I estimate separate CATEs for those whose number of days since recent purchase is above and below the median, respectively. Column (3) shows the p-value from a test of equal conditional average treatment effects for the two segments of customers, corresponding to that of β_3. Similarly, columns (4), (5), and (6) report the results on the propensity to share outcome.

Figure 1.4: Non-first-tier cities and first-tier cities CATE

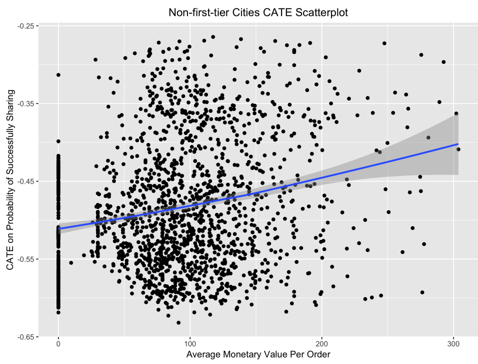

(a)

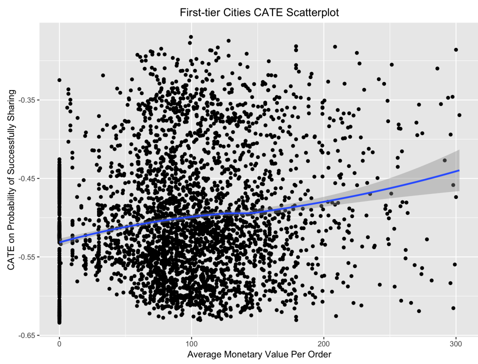

(b)

Moreover, I also find that those who had smaller average order size from past purchases tend to be more productive parents and have higher propensity to share (see Table 1.8 and Appendix A.4 Tables A.2 and A.4, and for other extensive margin outcomes see Appendix A.4 Tables A.5, A.6, and A.7). One might argue that, to some extent, the average order size might also reflect a customer's propensity for stockpiling more than one's purchase power and degree of price sensitivity; this, however, is not a concern here because the grocery partner mostly carries perishable products. Furthermore, this heterogeneity finding is also consistent with the city-level heterogeneity I presented in the last section. For both customers from first-tier cities and non-first-tier cities, higher average order size correlates with a lower propensity to share (see Figure 1.4). The scatterplots for the outcome number of customers introduced by average order size within first-tier and non-first-tier cities do not show any obvious patterns.

Knowledge of the Promoted Products Quality

I hypothesize that those customers who have consumed the promoted products before are less hesitant toward sharing and more likely to be productive parents. One reason is that these customers have more knowledge and therefore less uncertainty about the product quality. It could also be the case that they decide to include some positive comments when they share (although this information is not visible to the researcher). I compute two indicators of whether a customer has purchased the exact promoted bundle before and whether a customer has purchased just the promoted products before, regardless of the quantities in the bundle. The finding indeed supports the hypothesis (see Tables 1.8, and Appendix A.4 Tables A.2, A.3, and A.4; for other extensive margin outcomes see Appendix A.4 Tables A.5, A.6, and A.7).

Heterogeneity by Customer Gift Card Usage

I expect that those who have used more gift cards in their past purchases to be less productive parents and less willing to share. This is because, relatively speaking, these are more passive customers; they choose the platform because they have received gift cards. Specifically, I capture one's gift card usage with two measures, the average monetary value of gift card one applies per order among their purchases from the past year and the number of total orders that they use gift cards for in the past year. For brevity, I present results on the first measure here and refer the readers to the appendix for similar results using the second (see Table 1.8 and Appendix A.4 Table A.2).

Diminishing Marginal Return on Social Promotions - Recency, Frequency, and Degree of Customer Activity

As a growth strategy, social promotions may incur diminishing marginal return. Over time, customers who have been engaging in social promotions may generate fewer leads and have a declining propensity to share. One possible reason is sharing fatigue. They may be concerned that repetitively distributing similar promotional information will disturb friends.

To test this, given that the firm does not directly track customers' past responses to social promotions, I compute classical customer-value measures recency and frequency as proxy variables to a customer's likelihood of having engaged in social promotions before the experimental intervention. This assumes that those who bought more recently and more frequently are more likely to have used some social coupons. I find that the results indeed support my hypothesis. Among customers who have made at least one prior purchase, those whose recent purchase occurred longer ago and bought less frequently in the past tend to have higher propensity to share and lead-generation potential (see Table 1.8 and Appendix A.4 Tables A.3 and A.4; for other extensive margin outcomes see Appendix A.4 Tables A.2, A.5, A.6, and A.7).

Furthermore, active customers are more likely to have engaged in social promotions. Following the literature convention, these refer to those who have made at least one prior purchase in the past year. Despite that active and inactive customers do not exhibit significantly different lead-generation potential, inactive customers do have significantly higher propensity to share, 54% versus 49% (see Table 1.8).

These findings provide suggestive evidence that frequent social promotions can lead to sharing fatigue. A platform that wishes to continue utilizing social promotions to grow should consider experimenting with other designs. It may not be desirable to always inform a customer upfront about the precise products on promotion and the precise discount. As witnessed in the online markets, a firm may introduce some uncertainty to drive continued customer engagement. For example, the firm can make the customer only see the potential set of items that qualify for promotion, or just the range of potential discounts before she shares the deal. Upon successfully sharing, the firm reveals to her the entire deal, i.e. the exact product she can purchase at a discount and the precise depth of promotion. Alternatively, the firm can impose mandatory sharing only on a random or targeted subset of customers in each campaign; and each campaign entails a different segment of customers.

1.10 Conclusion

Firms continue to innovate on bundling promotions with social incentives. Yet standard field experimentation does not suffice to provide causal answers for measuring the social impacts thereof. This paper provides a first attempt to address the associated challenges when the social incentive is a sharing requirement. It contributes novel empirical strategies for conductitng valid inferences and provides some substantive strategic insights through an online field experiment. The proposed methods flexibly accommodate the amount of information available to the firms that engage in these practices. I find that social promotions can serve as an effective growth strategy for customer base activation. The findings are relevant for firms that wish to leverage the social channel to prosper.

This research has some important limitations; the results and methods might be best viewed as starting points for future research. The data contain a single experiment from a specific product market, hosted by one e-tailer site that is at a specific stage of growth. It would be valuable to replicate the study to see whether social promotions remain effective in other product categories. It is unclear if the findings on customer substitution would generalize. Moreover, as new customer acquisition and existing customer retention strategies, do social promotions suffer from diminishing marginal returns as a firm grows? Would social promotions work at all for companies that are already at the maturity stage and have achieved high market penetration?

This paper also identifies other exciting areas for future research. While this paper concerns the efficacy of social promotions in the short run by measuring immediate profitability, the long-run effects remain to be investigated. Prior literature suggests that in the context of commercial banking, referral customers are more valuable than other customers because they exhibit higher margins (which erode over time) and persistently lower churn [Schmitt et al., 2011]. It would be interesting to see whether similar results hold in the consumer packaged goods industry, i.e. whether referrals are more profitable and more loyal. In addition, social promotions can benefit a firm in at least two ways in the long run: they can improve brand awareness and accumulate additional data assets.

Furthermore, best practices of social promotions need to take many factors into account. As briefly discussed in section 1.9, best practices of social promotions need to take many factors into account. For instance, when running social promotions, should a firm use them on multiple products across different product categories? Can frequent social promotions lead to customer fatigue? Alternatively, can frequent social promotions decrease customers' sharing costs by habituating them into voluntary information sharing on the firm's behalf? Would other social incentives, such as randomized rewards for successfully sharing, be better?

Finally, I also do not address the competitive aspects of social promotions. How would social promotions affect consumer behavior and firm strategies when multiple firms compete for attention on social media? The answers to these questions are unclear ex-ante and yet would be of great interest both practically and theoretically. In my ongoing research, I attempt to address some of these questions. In my own humble opinion, social marketing is the wave of the future.

2

Optimal Targeting in Social Promotions

2.1 Introduction

This chapter examines the different targeting schemes a firm may wish to consider when deciding whether to impose a social incentive on top of a standard promotion, i.e. also requiring the customer to share the promotion with friends on social media to qualify for the discounts. A major challenge herein is whether or not to target the social incentive to a focal customer also affects how many leads will be generated. This paper proposes several different targeting schemes and compare their performances through counterfactual estimations.

Ideally, a firm may wish to mandate sharing only on customers who contribute net incremental benefit to the platform (with versus without the social incentive). These customers are the ones who incur relatively low sharing costs and yet exhibit high lead-generation potential. The presence of referrals presents a unique challenge for the targeting decision in our context.

2.2 Post-experiment Optimal Targeting

We start with a post-experiment targeting scheme which addresses the challenge herein with an "effective" profit idea. We find that compared with requiring all to share, requiring a subset of about 70% of the customers to share would increase the firm's profit by about 1%. The characteristics of these customers are consistent with the type of customers with high lead-generation potential.

Given that there is only one generation of children customers in our setting, it is straightforward to directly attribute children's profit contribution to the parent customers. We refer the readers to the first chapter for relevant definitions and notations. Formally, the optimal targeting policy maximizes the expected aggregate profit. In the context of social promotions, this objective function can be expressed as the sum of all organic customers' "effective" profit, i.e.

$$\begin{aligned} d^* &= \arg\max_{d \in \mathbb{D}} E[P^{\text{adj}}(d)], \\ &= \arg\max_{d \in \mathbb{D}} E[\sum_{i \in \Omega} \{P^{\text{adj}}_{i,} * (1 - d_i) + P^{\text{adj}}_{i,\text{NS}} * d_i\}], \end{aligned}$$

where $E[P^{\text{adj}}(d)] = E_{\mathbb{X}}[E[P^{\text{adj}}(d)|X_1,\ldots,X_N]]$ is the expected aggregate profit when implementing the targeting policy d on N organic customers who are randomly selected from the entire customer population, O is the set of organic customers, and X_i denotes customer i's set of observed characteristics.

Assume that the SUTVA holds among organic customers, maximizing the expected aggregate profit is equivalent to maximizing each organic customer's "effective" profit, namely

$$d_i^* = \arg\max_{d_i \in \{0,1\}} E[P_{i,}^{\text{adj}} * (1-d_i) + P_{i,\text{NS}}^{\text{adj}} * d_i | X_i = x].$$

Then, it follows that $d_i^* = 1$ if and only if $E[P_{i,\text{NS}}^{\text{adj}} | X_i = x] > E[P_{i,}^{\text{adj}} | X_i = x]$, and $d_i^* = 0$ otherwise. As a reminder, $d_i^* = 1$ implies that the optimal policy d^* assigns customer i to NS. Substituting in the definitions of the conditional average treatment effects (CATE) τ_i^{ExtP} and τ_i^{IntP}, we have

$$\begin{gathered} E[P_{i,\text{NS}}^{\text{adj}}] > E[P_{i,}^{\text{adj}}], \\ \Longleftrightarrow \\ (1-z_i)(P_i^{\text{adj}} + \tau_i^{\text{ExtP}} + \tau_i^{\text{IntP}}) + z_i P_i^{\text{adj}} > (1-z_i)P_i^{\text{adj}} + z_i(P_i^{\text{adj}} - \tau_i^{\text{ExtP}} - \tau_i^{\text{IntP}}), \\ \Longleftrightarrow \\ \tau_i^{\text{ExtP}} + \tau_i^{\text{IntP}} > 0. \end{gathered}$$

Intuitively, a firm should assign customer i to NS if and only if her net incremental profit $\tau_i^{\text{ExtP}} + \tau_i^{\text{IntP}}$ is positive when not required to share. In other words, the social obligation is unsuitable for a customer when the potential profit loss from herself under mandatory sharing outweighs the potential profit gains she would contribute from her referrals.

Evaluation of Targeting Policies

This subsection presents the unbiased estimator I propose to evaluate a counterfactual targeting policy's aggregate profitability. The objective function can be expressed as the aggregate "effective" profit from all organic customers, i.e.

$$\begin{aligned} \widehat{P^{\text{adj}}(d)} &= \sum_{i \in \Omega} [\widehat{P_{i,}^{\text{adj}}} * (1 - d(X_i)) + \widehat{P_{i,\text{NS}}^{\text{adj}}} * d(X_i)] \\ &= \sum_{i \in \Omega} \Big\{ (1-d_i)[(1-z_i)P_i^{\text{adj}} + z_i(P_i^{\text{adj}} - \hat{\tau}_i^{\text{ExtP}} - \hat{\tau}_i^{\text{IntP}})] \\ &\qquad + d_i[(1-z_i)(P_i^{\text{adj}} + \hat{\tau}_i^{\text{ExtP}} + \hat{\tau}_i^{\text{IntP}}) + z_i P_i^{\text{adj}}] \Big\}, \end{aligned}$$

where P_i^{adj} denotes customer i's observed "effective" profit, $\hat{\tau}_i^{\text{IntP}}$ and $\hat{\tau}_i^{\text{ExtP}}$ denote customer i's conditional average treatment effects measured in terms of profitability. Therefore, calibrating profit of a targeting policy requires estimating each customer's conditional average treatment effects (CATE). I assume that the observed characteristics X_i captures the type of customer i, and obtain the CATE estimates from the grf model.

One can show that

$$E[\widehat{P^{\text{adj}}(d)}] = E[P^{\text{adj}}(d)].$$

Results and Discussion

I find that compared with requiring all to share, requiring a subset of about 70% of the customers to share would increase the firm's profit by about 1%. The profitability improvement under optimal targeting is highly statistically significant compared with forcing a random half of the customers to share, and marginally significant compared with forcing all and a random 70% of the customers to share (see Tables 2.1 and 2.4). Appendix B.2 Profit Figure B.10 presents the optimal intensity-profit pair and the 95% confidence region using the standard bootstrap technique. Appendix B.2 Profit Figure B.11 presents the corresponding information on percents of increase. For a more global perspective on how optimal targeting performs compared with random targeting, see Appendix B.2 Figures B.12 and B.13.

Table 2.1: Summary statistics

	Average "effective" Profit (USD)	Standard Error (USD)	Percent of Increase
Optimal (32.2 %NS)	0.4159	0.0475	1.59%
0%NS	0.4108	0.0497	0.29%
Random 30% NS	0.4106	0.0497	0.25%
Benchmark (50% NS)	0.4096	0.0497	-

Note: This table presents summary statistics on how the proposed targeting scheme compares with exposing nobody and a random 30% of customers to standard promotion in terms of the "effective" profit for organic customer. Standard errors are computed using bootstrap. The comparison benchmark is the realized social promotion, where a random half of the customers are obliged to share.

Table 2.2: Pairwise t-test between counterfactual policies

	0%NS	Random 30% NS	Benchmark (50%NS)
Optimal (32.2%NS)	0.089 (1.699)	0.056 (1.913)	0.007*** (2.693)
0%NS	-	0.832 (0.212)	0.325 (0.985)
Random 30% NS		-	0.439 (0.773)

Notes: This table reports the p-value with t-statistic for pairwise comparison of different targeting policies.

Moreover, the characteristics of the counterfactual S customers are consistent with the type of customers with high lead-generation potential. Under counterfactual optimal targeting, compared with customers assigned to regular coupons, those assigned to social coupons, on average, have purchased less frequently in the past year, have been registered for a longer period of time, have been more price sensitive, and have used fewer gift cards (see Table 2.3 for covariates comparisons among the counterfactual S and NS customers under optimal targeting and Appendix B.1 Figures B.1, B.2, B.3, B.4, B.5, B.6, B.7, and B.8 for more details on comparisons of their cumulative distribution functions). The counterfactual S customers also used direct discounts (the major type of alternative promotional vehicle to coupons that do not require redemption) less frequently in the past year. Furthermore, the fraction of new customers is also smaller among those assigned to the social coupon versus those assigned to the regular coupon; this suggests that social promotions may work better on existing customers than new customers. (see Appendix B.1 Figure B.9). The only dimension inconsistent with the observed heterogeneity results is that optimal targeting suggests imposing the social incentive on a relatively small fraction of customers who have consumed the promoted products before. This, however, is sensible, given that grf ranks whether the customer has consumed the promoted products before relatively unimportant (see Appendix A.4 Table A.3).

Table 2.3: Covariates comparison for optimal targeting within 50% NS cities

	Both	Target NS Customers	Target S Customers	t.p	u.p	ks.p
Past Purchase Behavior						
Days Since Recent Purchase	59.182	43.103	62.904	0	0	0
Number of Orders Last Year	18.677	34.219	15.333	0	0	0
Purchased Exact Promoted Bundle Before	0.332	0.451	0.307	0	0	0
Purchased Promoted Products Before	0.381	0.514	0.353	0	0	0
Experiment Specific						
New Customer	0.098	0.132	0.082	0	0	0.003
Days Since Registration	415.798	288.064	443.278	0	0	0
Referred by Others	0.036	0.205	0	0	0	0
Response to Past Marketing Activities						
Number of Orders Using Gift Cards Last Year	2.244	6.166	1.400	0	0.00000	0.0001
Number of Orders with Direct Discount Last Year	8.830	13.975	7.723	0	0.00000	0
Average Monetary Value of Coupons Per Order	34.629	16.346	38.563	0	0	0
Average Monetary Value of Gift Cards Per Order	10.939	17.928	9.436	0.00000	0.00002	0.003
Average Monetary Value of Coupons and Direct Discount Per Order	43.960	25.095	48.019	0	0	0
Number of observations	5,959	1,916	4,043			

Notes: This table presents the set of observable customer characteristics that differ significantly between the counterfactual NS and counterfactual S customers under optimal targeting. Columns (1), (2), and (3) report the means of individual characteristics in both counterfactual NS and counterfactual S, counterfactual NS, and counterfactual S subsamples, respectively. Columns (4), (5), and (6) report the p-values of t-test, Mann-Whitney U-test, and Kolmogorov-Smirnov test on whether counterfactual S and counterfactual NS are significantly different in terms of the mean, median, and empirical distributions of each covariate, respectively.

2.3 Optimal Targeting using Contextual Bandits

While the previous section proposes a post-experiment targeting scheme based on the “effective” profit idea using grf CATE estimates, this section proposes an alternative that taps into incoming data real time to balance exploration and exploitation.

One issue with the approach in the prior section is that running the A/B test posts an opportunity cost on the firm [Feit and Berman, 2019]. Moreover, in a fast-evolving social e-commerce landscape, ideally, a firm may wish to tap into incoming data assets real time to inform its targeting decisions. This section proposes an alternative of applying contextual bandits instead of ex-post targeting based on standard A/B testing. Contextual bandits use data more efficiently by balancing exploration and exploitation as each customer arrives [Scott, 2010]. In contrast, A/B testing can be viewed as an exploration-first greedy algorithm, where the data-collection stage is mere exploration and the targeting stage is mere exploitation.

Furthermore, what distinguishes our targeting setting from a standard one is the presence of referrals. Whether or not these referrals arrive in the first place can depend on the targeting assignment of other customers who have already been exposed to the promotion.

As a starting point, to assess how the bandits perform in comparison, we do not distinguish between whether a customer is a referral or not. We directly apply the contextual bandits in offline simulations. We present simulation results using two popular contextual bandits, linear Thompson sampling (LinTS, [Agrawal and Goyal, 2012]) and linear upper confidence bound (LinUCB, [Lihong et al., 2010]). Our simulations use the experimental data. To evaluate the performance of each algorithm, we use grf CATE estimates and a doubly robust estimator.

Figure 2.1 shows that LinUCB and LinTS perform similarly. On average, both suggest targeting the social incentive to about 45% of customers. They are significantly better than the random targeting benchmark (see Table 2.4). While LinUCB has a smaller point estimate, it also has a smaller standar error. Figure 2.2 presents the average fraction of customers assigned to S or NS respectively at each step in the simulation.

Figure 2.1: Comparison of Contextual Bandits v. Random Targeting

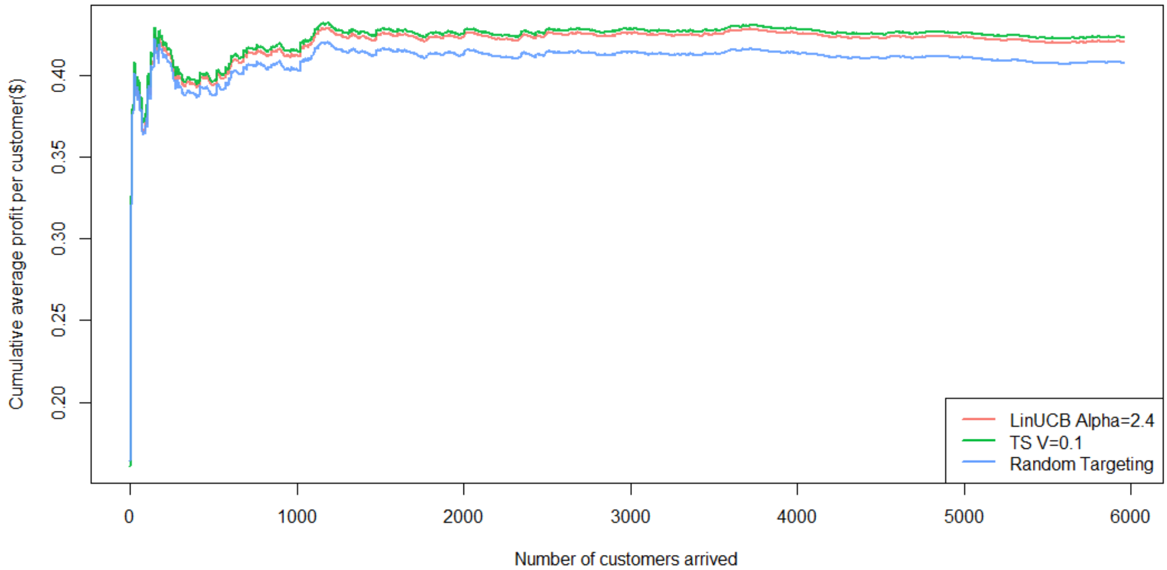

Note: This plot shows the cumulative average profit per customer under the doubly-robust estimator as each customer enters the simulation [Dudik et al., 2011]. For each of the three different targeting schemes, we present how the point estimates evolve as more customers enter the promotion. The top two lines present the average estimates of running the simulation 100 rounds using LinTS and LinUCB, respectively. We tune the hyperparameters α and V using the first 1000 customers entering the simulation. The bottom line presents our comparison benchmark of random targeting, i.e. assigning customers in standard A/B test.

Table 2.4: Summary statistics

	Average Profit per Customer(USD)	Standard Error (USD)	Percent of Increase
LinUCB (55% NS)	0.420	0.008	3.2%
LinTS (55% NS)	0.423	0.013	3.9%
Benchmark (50% NS)	0.408	0.001	-

*Notes:*This table reports how the contextual bandits perform compared with the random targeting benchmark.

Figure 2.2: Customer assignment under LinUCB

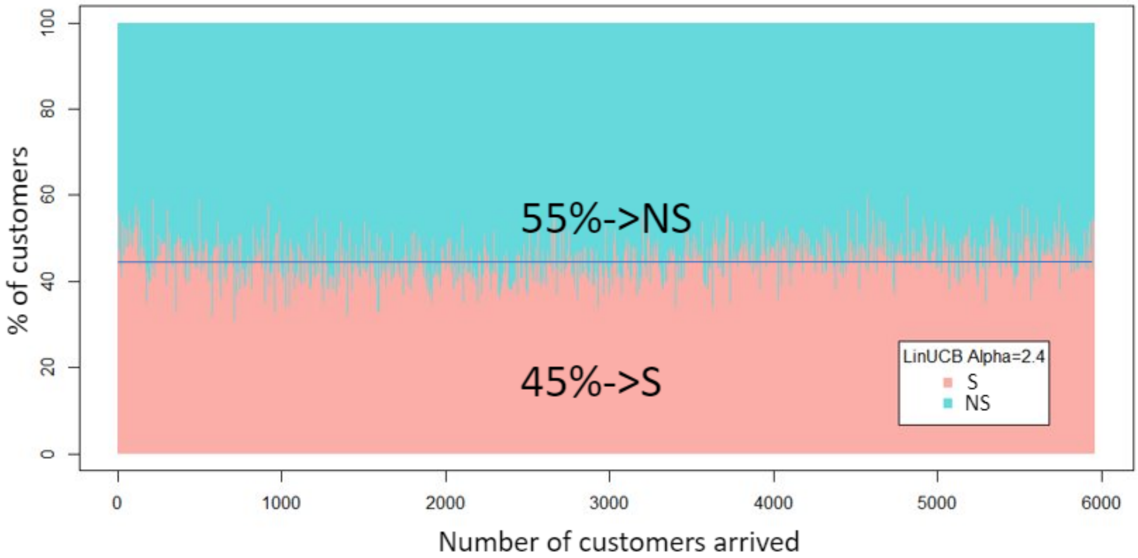

Note: This plot shows the average fraction of customers assigned to either S and NS at each step in our simulation.

2.4 Limitations

The proposed targeting schemes have several critical limitations. First, none of the proposed targeting schemes both directly address the presence of referrals and also use incoming data real time. We conjecture that we can synthesize the "effective" profit idea with the contextual bandits. By recursively attributing referral profits to corresponding parent customers, the contextual bandit will prescribe an action for each customer based on her incremental "effective" profit contribution. More precisely, for each arriving customer, we can construct the "effective" confidence interval to incorporate referral profits. Meanwhile, we will also update the CATE estimates as the contextual bandits are running. We plan to run additional simulations to compare the performance and then run online experiments with the collaborating firm to further test the performance.

While we have shown that the presence of referrals in our context does not make offline counterfactual evaluation problematic, we conjecture that the proposed schemes may not suffice when the referrals account for a relatively large fraction of arriving customers. We conjecture, however, the "effective" LinUCB and "effective" LinTS can overcome the difficulties.

Furthermore, we do not directly consider peer effects among organic customers. However, we conjecture that under the two-level randomization design, the alternative unbiased estimator (proposed for predicting expected counterfactual targeting profits) would remain valid and allow for direct incorporation of the peer effects. The idea of attributing children's profit to parents in our proposed framework should also generalize when multiple generations of referrals are present.

3

Bundling in E-commerce - an Empirical Case Study on the DSLR Camera Market

3.1 Introduction

Bundling is a prevalent business strategy in e-commerce. On Taobao.com, the world's largest online marketplace, electronic products sellers frequently bundle a big-ticket item (such as a smart-phone, camera, or game console) with a variety of accessories (such as a memory card, a case, and a screen protector). When a customer searches for a base good, such as a DSLR camera, she frequently finds herself confronted with multiple bundles to choose from in addition to the base good itself.[1] Other leading e-commerce platforms, such as Amazon and B&H, also offer similar bundles. This kind of bundling is known as mixed bundling in the literature because both the bundle and each individual product component in the bundle are available for purchase. Yet there has been scant empirical studies on the pricing and design of mixed bundling in online markets (see Venkatesh and Mahajan [2009] for a comprehensive survey of bundling in the marketing literature).

Conventional wisdom suggests that in mixed bundling, a bundle should cost less than its different parts combined. Since Stigler [1963], there has been substantial theoretical economics research on when bundling is a profitable strategy and how bundling affects consumer welfare and overall market efficiency. However, it remains a theoretical challenge to characterize equilibrium pricing of mixed bundling in an oligopoly context. This paper aims to fill in the corresponding empirical gap by studying the design and pricing of bundling in online markets.

Unlike what conventional wisdom suggests, we have found that on Taobao.com, a bundle tends to cost even more than purchasing each bundle item separately from the same online market. This "bundle premium" phenomenon permeates multiple product categories, including DSLR cameras, iPads, cell-phones, and video game consoles, just to name a few. Even conservatively speaking, these "premiumed bundles" account for 30% of market share. Furthermore, the magnitude of the premiums (e.g. an average of 17 USD, or 4% of the primary base good, for Canon 700D) is also beyond that of rational expectation. Search cost alone does not suffice to account for the magnitude of the premium, as an average customer's time value for running a search for each accessory on Taobao.com is way lower. Admittedly, it still is mentally burden-

[1]Taobao typically requires a customer to make an active choice among the different bundles and the stand-alone base good. Each seller may list up to eight bundles in addition to the base good itself.

some for customers to compare the different listings for each accessory given the substantial quality uncertainty on Taobao. An alternative explanation is that of framing effect: when consumers are comparing cameras worth of hundreds of dollars, an extra ten dollars do not seem much. In other words, consumer's price sensitivity to these add-ons is fairly low. Moreover, transaction records show that the majority of customers choose to buy some bundles instead of just the base good (65% for Canon 700D).

The pilot study in this chapter focuses on bundles of one popular base good, the Canon 700D DSLR camera, as this camera enjoys the largest sales within the category and DSLR cameras generally have a rich set of add-on accessories. We find that Canon 700D buyers are paying, on average, \$5 to \$20 more for bundles compared to buying the base good from the same seller and the exact same accessories from comparable Taobao sellers. The premium characterization is robust to several ways of measuring accessory prices (e.g. bootstrapping from the set of outside accessories listings), as well as controlling for seller reputation.

3.2 Data and Summary Statistics

Taobao.com provides a rich setting to study the design and pricing of bundling in online markets. We collect our data in the following manner. For example, for bundles of the base good Canon 700D, we search for the keyword "Canon 700D" on Taobao.com and collect all related listings. The information includes the bundle design, pricing, transaction records (recent 90 days), and customer review information. For each accessory, we run a keyword search on Taobao.com, and collect all matching listings on the first search results page.

First, we find that in the Canon 700D market, bundles account for 65% of purchases among all Canon 700D related transactions. Only 35% of customers decided to only purchase the base good Canon 700D alone. In general, a bundle with a higher index (displayed in a lower position on the item page) tends to have lower market share. The distribution of sales relative to bundle index is also presented in box plots and we observe some variation. Moreover, there is also substantial variation in pricing but generally each seller orders the listing of bundles in a price-increasing manner.

For the ease of exposition, we first introduce some notations before discussing how to quantify the bundle premiums. Let $p(B_j)$ denote the price of bundle j (where $1 \leq j \leq 8$, and B_0 is just the base good). From now on, we will refer to B_1 as the baseline bundle.

Let p_{acc} denote the "street price" for an accessory, i.e. the price a customer will pay if she were to purchase the accessory outside of the current store and on the street. To quantify the bundle premiums, we start with a relatively conservative measure of the "street price," the mode price of all matching accessory listings on the first search result page on the same e-commerce platform Taobao.com as our benchmark "street price." We subsequently bootstrap from these "street" accessories to show robustness of our results.

Let Δp denote the Bundle Premium of B_j rel to B_k,

$$\Delta p = p(B_j) - p(B_k) - \sum_{\forall \text{acc} \in B_j \& \notin B_k} p_{\text{acc}}.$$

Let Δn denote the # of accessories in B_j and not in B_k. To characterize the premiums, we begin by using the baseline bundle as the reference benchmark. Compared with using the base good

Figure 3.1: Summary Statistics for 700D Pilot Study

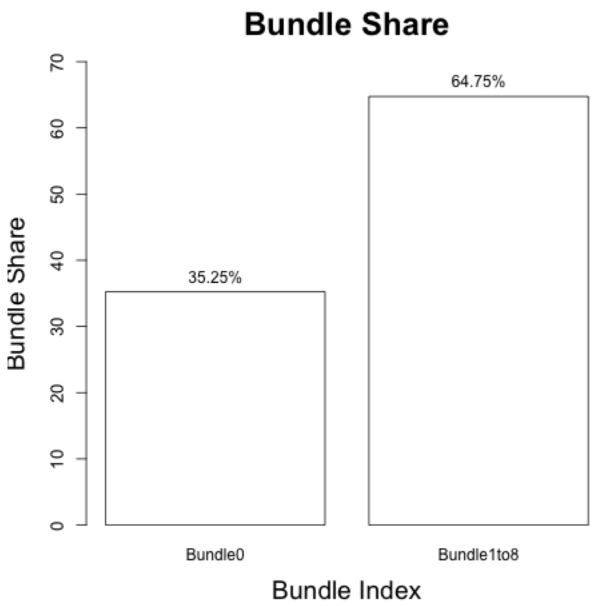

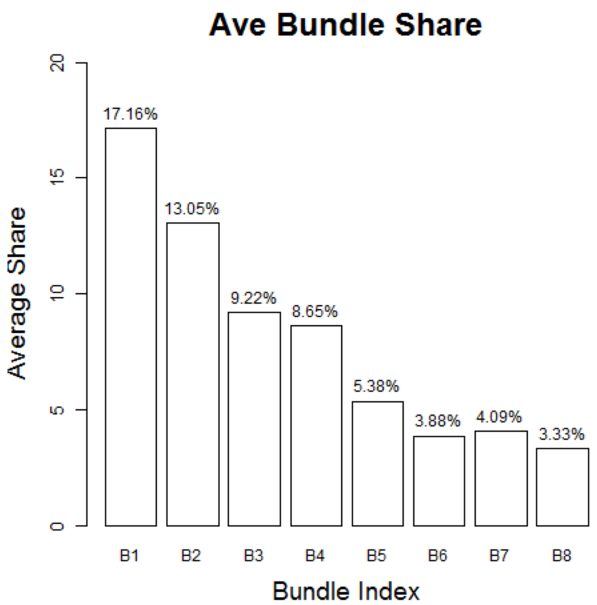

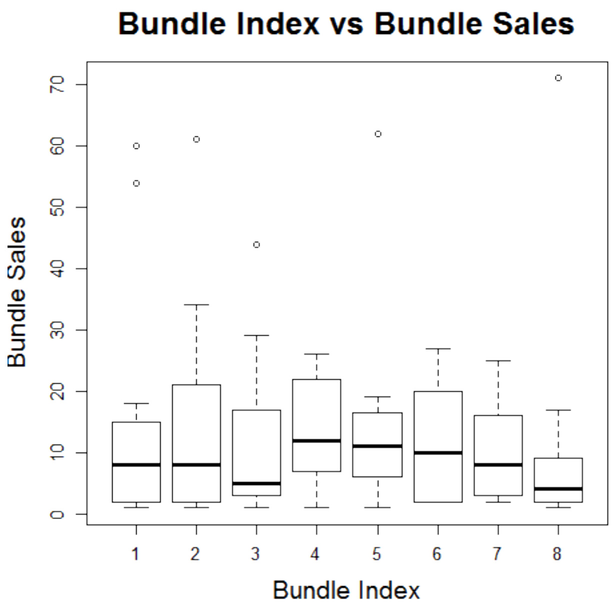
Bundle Index vs Bundle Sales
Bundle Sales
0
10
20
30
40
50
60
70
1
2
3
4
5
6
7
8
Bundle Index

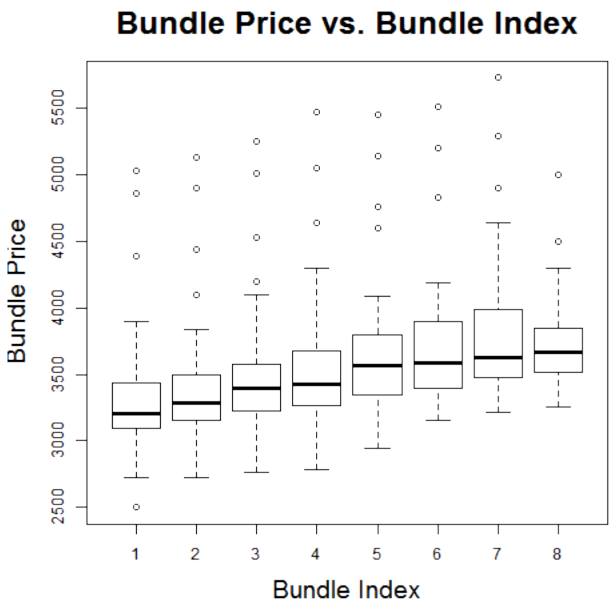
Bundle Price vs. Bundle Index
Bundle Price
2500
3000
3500
4000
4500
5000
5500
1
2
3
4
5
6
7
8
Bundle Index

Figure 3.2: Summary Statistics for 700D Pilot Study

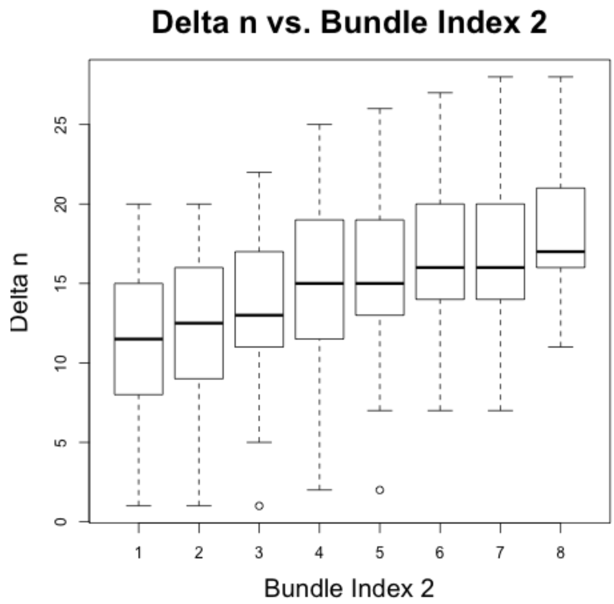

Figure 3.3: Summary Statistics for 700D Pilot Study

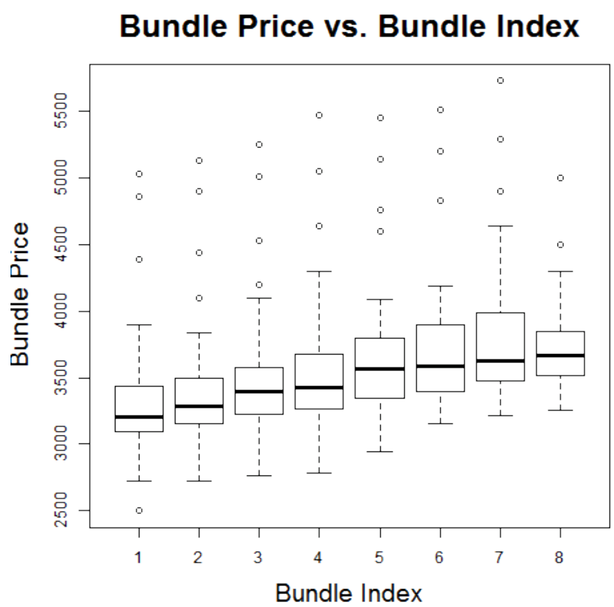

alone instead, this baseline bundle benchmark reduces measurement error as even the baseline bundle has on average 12 accessories. We present the findings in the figures below.

3.3 Models and Results

We answer the following research question: Suppose a seller observes a collection of bundles $\mathcal{B}$ listed in the marketplace and needs to design in his/her own store what bundles to offer and what price to charge for each bundle position (or bundle index j).

We make the following assumptions. Sellers in the marketplace follow some market norm to design the baseline bundle B_1 (both its content and pricing) and the number of accessories to include for each bundle position.

We highlight some preliminary reduced-form results based on our Canon 700D pilot study. First, we observe that among the accessories, memory cards, UV filters, and camera bags appear most frequently. Within a seller, these accessories are also most frequently upgraded in a higher position bundle compared with the baseline bundle. We run the following seller-fixed effect models.

$$\begin{aligned}\Delta p = \alpha\Delta n + \beta\text{bundle\_index} \\ +\gamma_1\mathbb{1}(\text{memory card (cc)} \in \Delta n) \\ +\gamma_2\mathbb{1}(\text{uv filter (uv)} \in \Delta n) \\ +\gamma_3\mathbb{1}(\text{camera bag (xb)} \in \Delta n)\end{aligned}$$

Table 3.1 reports the regression results. Controlling for seller fixed effects, on average a seller should price each higher bundle position with an incremental premium of 42 to 43 RMB. Each memory card contributes a premium of 35 to 41 RMB, each UV filter contributes a premium of 89 to 92 RMB and each camera bag contributes a premium of 70 RMB.

3.4 Conclusion

To conclude, our Canon 700D pilot study suggests that a prominent fraction of e-commerce practices contradicts conventional wisdom on bundling. Firms glean a bundle premium from consumers. Such practices are similar in nature to other firm shrouding behavior of negative product attributes. Specifically, firms shroud the add-ons without supplying consumers with any reference prices for the add-ons. Unless a consumer is diligent enough to conduct the search herself, she will just be ripped off by the seller and bear an extra markup that is beyond her rational convenience cost.

Future research may wish to structurally separate the different drivers of the bundle premiums to provide insights into the impact of such practices on consumer welfare and its implications for platform design and governance.

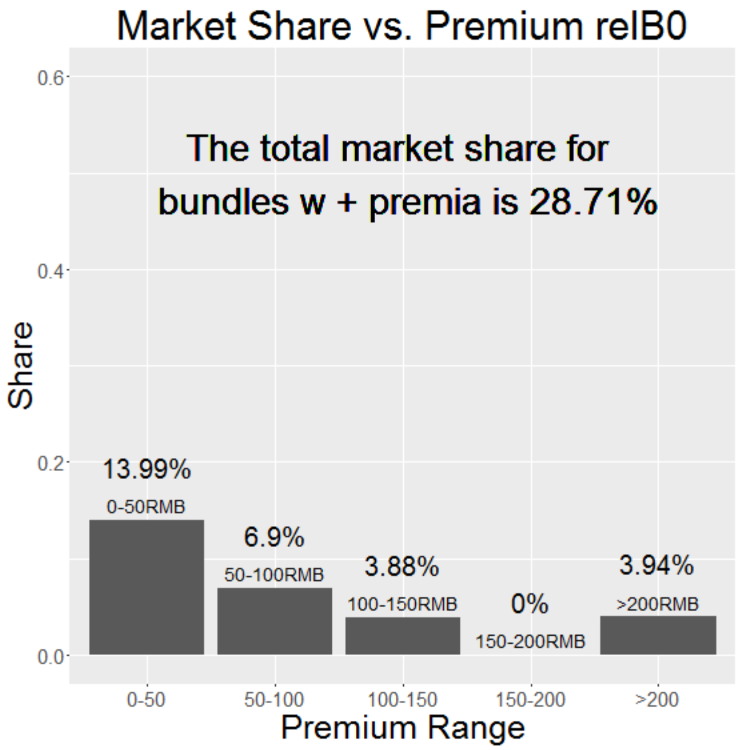

Figure 3.4: Summary Statistics for 700D Pilot Study

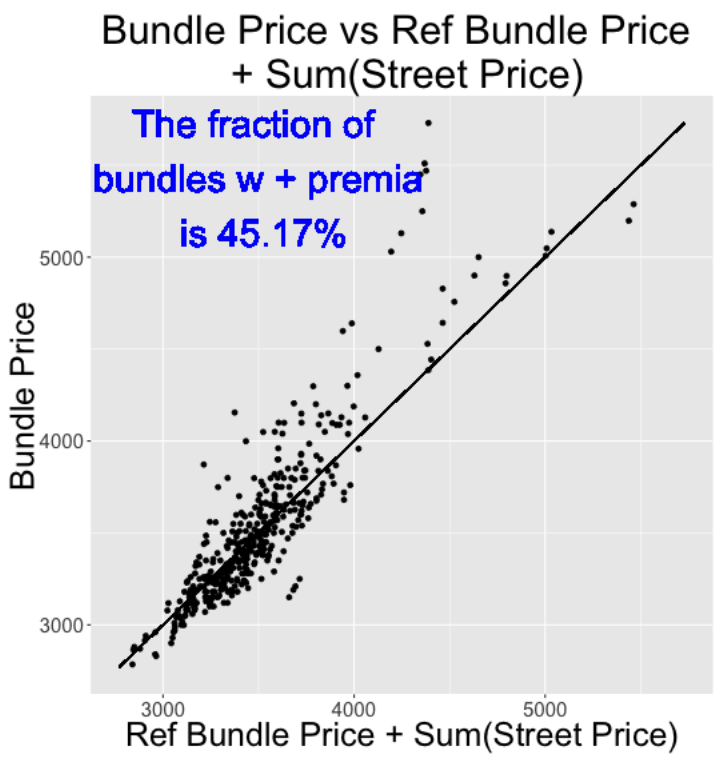

Figure 3.5: Summary Statistics for 700D Pilot Study

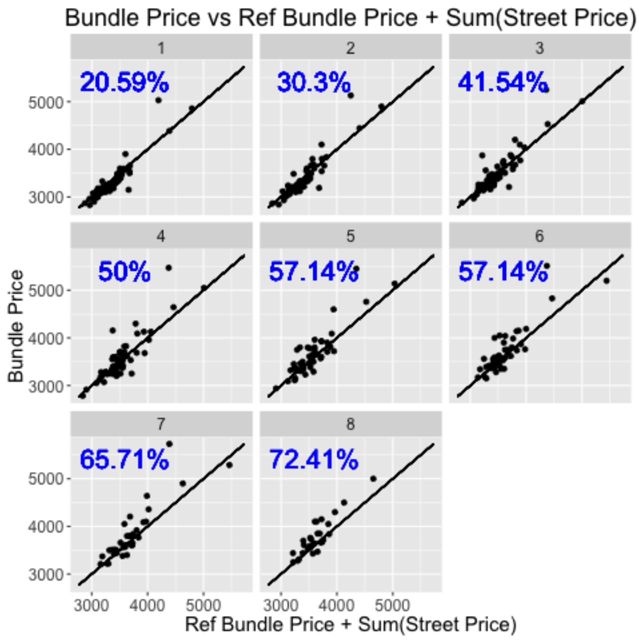

Figure 3.6: Summary Statistics for 700D Pilot Study

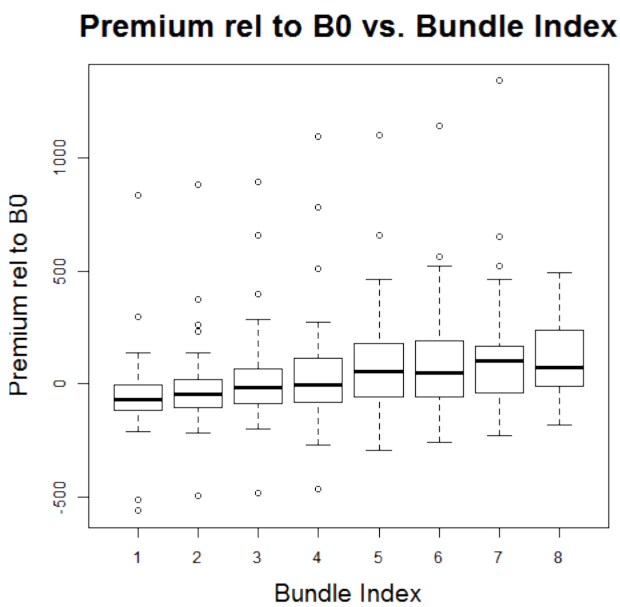

Figure 3.7: Summary Statistics for 700D Pilot Study

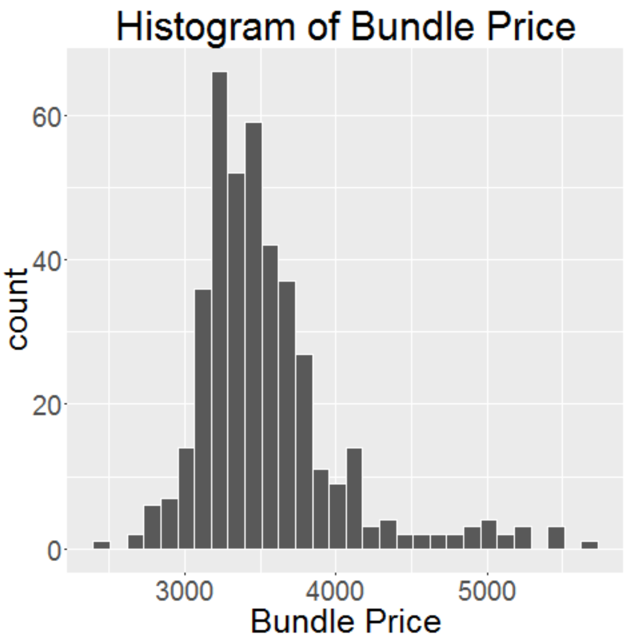

VARIABLES	(1) premium	(2) premium
number of different accessories relative to baseline bundle	-12.04**	-17.53***
	(5.701)	(6.413)
bundle_index2	41.67***	43.17***
	(4.671)	(4.798)
cc	34.86**	40.85***
	(13.85)	(13.90)
uv	88.87*	92.27**
	(45.91)	(45.18)
xb		70.41
		(47.25)
Constant	-58.23***	-60.78***
	(21.47)	(21.99)
Observations	352	352
R-squared	0.338	0.354
Number of item_id	67	67

Robust standard errors in parentheses
*** $p<0.01$, ** $p<0.05$, * $p<0.1$

Table 3.1: Regression Results

Bibliography

Shipra Agrawal and Navin Goyal. Thompson sampling for contextual bandits with linear payoffs. *PMLR*, 28(3), 2012.

Joseph G Altonji, Todd E Elder, and Christopher R Taber. Selection on observed and unobserved variables: Assessing the effectiveness of catholic schools. *Journal of Political Economy*, 113(1):151–184, 2005.

Eva Ascarza, Peter Ebbes, Oded Netzer, and Matthew Danielsom. Beyond the target customer: Social effects of customer relationship management campaigns. *Journal of Marketing Research*, 54(3):347–363, 2017.

Susan Athey, Julie Tibshirani, and Stefan Wager. Generalized random forests. *arXiv preprint arXiv:1610.01271*, 2016.

Kyle Bagwell and Michael Riordan. High and declining prices signal product quality. *American Economic Review*, 81(1):224–39, 1991.

Eyal Biyalogorsky, Eitan Gerstner, and Barak Libai. Customer referral management: Optimal reward programs. *Marketing Science*, 20(1):82–95, 2001.

Miriam Bruhn and David McKenzie. In pursuit of balance: Randomization in practice in development field experiments. *American economic journal: applied economics*, 1(4):200–232, 2009.

Leonardo Burszytn, Florian Ederer, Bruno Ferman, and Noam Yutchman. Understanding mechanisms underlying peer effects: Evidence from a field experiment on financial decisions. *Econometrica*, July 2014.

Inyoung Chae, Andrew T Stephen, Yakov Bart, and Dai Yao. Spillover effects in seeded word-of-mouth marketing campaigns. *Marketing Science*, 36(1):89–104, 2016.

Xiaolong Chen. China city statistical yearbook. *China City Statistical Yearbook*, 2015.

Miroslav Dudik, John Langford, and Lihong Li. Doubly robust policy evaluation and learning, 2011.

Esther Duflo and Emmanuel Saez. The role of information and social interactions in retirement plan decisions: Evidence from a randomized experiment. *The Quarterly Journal of Economics*, 118(3):815–842, 2003.

Barzegary Ebrahim, Yoganarasimhan Hema, and Pani Abhishek. Personalized free trials: Design and evaluation.

Dean Eckles, Brian Karrer, and Johan Ugander. Design and analysis of experiments in networks: Reducing bias from interference. *Journal of Causal Inference*, 2017.

Pater Elstrom and David Ramli. Ex-google engineer builds $1.5 billion startup in 21 months., 2017.

Elea McDonnell Feit and Ron Berman. Test & roll: Profit-maximizing a/b tests. *SSRN*, 2019.

David Godes and Dina Mayzlin. Firm-created word-of-mouth communication: Evidence from a field test. *Marketing science*, 28(4):721–739, 2009.

Miguel Godinho de Matos, Pedro Ferreira, and Rodrigo Belo. Target the ego or target the group: Evidence from a randomized experiment in pro-active churn management. 2016.

Tong Guo, S Sriram, and Puneet Manchanda. The effect of information disclosure on industry payments to physicians. *Available at SSRN 3064769*, 2017.

Wesley R Hartmann. Demand estimation with social interactions and the implications for targeted marketing. *Marketing science*, 29(4):585–601, 2010.

Michael G Hudgens and M Elizabeth Halloran. Toward causal inference with interference. *Journal of the American Statistical Association*, 103(482):832–842, 2008.

Adobe Digital Insights. Adi – state of digital advertising 2018, 3 2018.

V Kumar, J Petersen, and Robert P. Leone. Driving profitability by encouraging customer referrals: Who, when, and how. *Journal of Marketing - J MARKETING*, 74:1–17, 09 2010. doi: SS‘10.1509/jmkg.74.5.1.

Krista J. Li and Sanjay Jain. Behavior-based pricing: An analysis of the impact of peer-induced fairness. *Management Science*, 62(9):2705–2721, 2015. doi: SS‘10.1287/mnsc.2015.2265.

Li Lihong, Chu Wei, Langford John, and Robert E. Schapire. A contextual-bandit approach to personalized news article recommendation. *WWW*, 2010. URL SS‘`http://rob.schapire.net/papers/www10.pdf`.

Puneet Manchanda, Grant Packard, and Adithya Pattabhiramaiah. Social dollars: The economic impact of customer participation in a firm-sponsored online customer community. *Marketing Science*, 34(3):367–387, 2015.

Charles F Manski. Identification of treatment response with social interactions. *The Econometrics Journal*, 16(1):S1–S23, 2013.

Mobile Marketer. Online retail is booming in China, 14 2019.

Edward Miguel and Michael Kremer. Worms: identifying impacts on education and health in the presence of treatment externalities. *Econometrica*, 72(1):159–217, 2004.

Paul Milgrom and John Roberts. Limit pricing and entry under incomplete information: An equilibrium analysis. *Econometrica*, 50(2):443–59, 1982.

Bang Nguyen and Lyndon Simkin. The dark side of CRM: advantaged and disadvantaged customers. *Journal of Consumer Marketing*, 30(1):17–30, 2013. doi: SS‘10.1108/07363761311290812.

Navdeep S Sahni, Dan Zou, and Pradeep K Chintagunta. Do targeted discount offers serve as advertising? evidence from 70 field experiments. *Management Science*, 63(8):2688–2705, 2016.

Philipp Schmitt, Bernd Skiera, and Christoph Van den Bulte. Referral programs and customer value. *Journal of Marketing*, 75:46–59, 01 2011. doi: SS‘10.2307/25764294.

Steven L. Scott. A modern bayesian look at the multi-armed bandit. *Applied Stochastic Models in Business and Industry*, 26:639–658, 2010. URL SS‘`http://onlinelibrary.wiley.com/doi/10.1002/asmb.874/abstract`.

Avni Shah, Matthew Osborne, Jaclyn Lefkowitz, Alissa Fishbane, and Dilip Soman. Can making family salient increase financial savings? quantifying heterogeneous treatment effects in voluntary retirement contributions using a field experiment in mexico. *SSRN*, 2019.

Betsy Sinclair, Margaret McConnell, and Donald P Green. Detecting spillover effects: Design and analysis of multilevel experiments. *American Journal of Political Science*, 56(4):1055–1069, 2012.

George J Stigler. United states v. loew's inc.: A note on block-booking. *Sup. Ct. Rev.*, page 152, 1963.

Christophe Van Den Bulte, Emanuel Bayer, Bernd Skiera, and Philipp Schmitt. How customer referral programs turn social capital into economic capital. *Journal of Marketing Research*, 55(1):132–146, 2018.

R Venkatesh and Vijay Mahajan. 11 the design and pricing of bundles: a review of normative guidelines and practical approaches. *Handbook of pricing research in marketing*, 232, 2009.

Appendix A

Appendix for Chapter 1

A.1 Robustness Checks on Selection on Unobservables

I apply the method developed in Altonji et al. [2005] and show that the normalized shift in the distribution of unobservables need to be fairly large relative to the shift in the observables in order to completely explain away the treatment effects of interest. Specifically, the assumption is that the standardized selection on unobservables equals that on observables, i.e.

$$\frac{E(\varepsilon|Z=1)-E(\varepsilon|Z=0)}{Var(\varepsilon)}=\frac{E(X^{\top}\gamma|Z=1)-E(X^{\top}\gamma|Z=0)}{Var(X^{\top}\gamma)}.$$

Then, for example, the ratio of normalized shift in the distribution of unobservables relative to that of observables on the intensive margin parameter is almost 9, i.e.

$$\frac{\hat{\text{Int}}^{50\%}}{Var(Z)/Var(\tilde{Z})*(E(\varepsilon|Z=1)-E(\varepsilon|Z=0))}=\frac{0.057}{0.0065}=8.8,$$

where $\tilde{Z}$ represents the residuals from the regression $Z=X^{\top}\gamma+\tilde{Z}$, and the denominator is the bias term for the true parameter Int in probability limit.

A.2 Computational Details on the Delta Method

This section presents the algebraic details for deriving the incremental propensity to buy for a 50% NS city vs. a 0% NS city as a function of the other causal parameters:

$$\Delta r_e(\alpha_1=50\%,\alpha_2=0\%)=r_e(\alpha_1=50\%)-r_e(\alpha_2=0\%)$$

$$=(1+\frac{1}{2}s(50\%))\frac{1}{2}(2r_S(50\%)+i(50\%))-(1+s(100\%))r_S(0\%)$$

$$=\frac{1}{2}(1+\frac{1}{2}s)(2r_S+i(50\%))-(1+s)r_S$$

The second step uses the assumptions that

$$r_S(50\%)=P_S(0\%)$$

and

$$s(50\%) = s(100\%)$$

for simplification. These assumptions can be relaxed as necessary.

Additionally, I obtain the heteroscedasticity consistent variance-covariance matrix for the parameter

$$\vec{\beta}^{\alpha,\gamma}$$

by jointly estimating the following equations

$$Y_i^I = r_S(0\%) + i(\alpha)Z_i + \varepsilon_i^I,$$

and

$$Y_i^E = s(1-\alpha = 100\%) - \text{Ext}^\alpha Z_i + \varepsilon_i^E.$$

Appendix Figures - Flow Charts

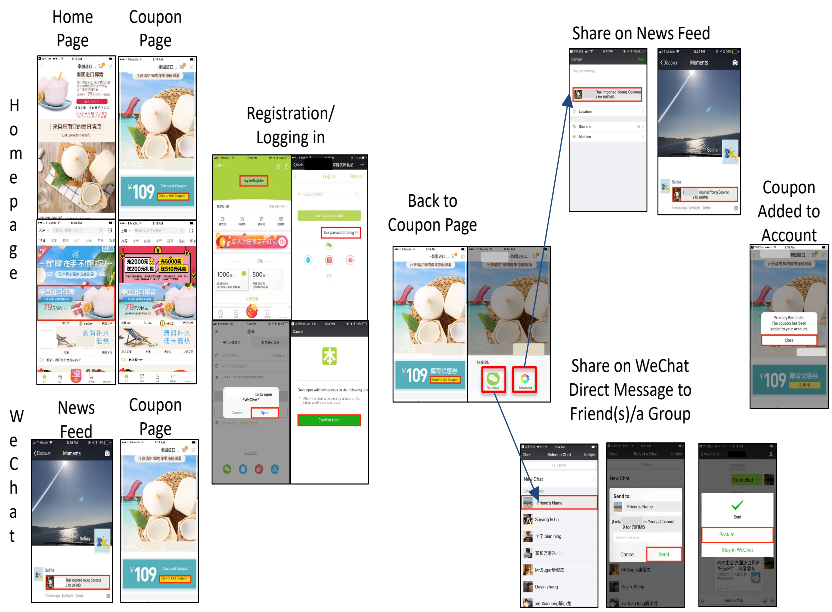

Figure A.1: Control City Coupon Activation Steps

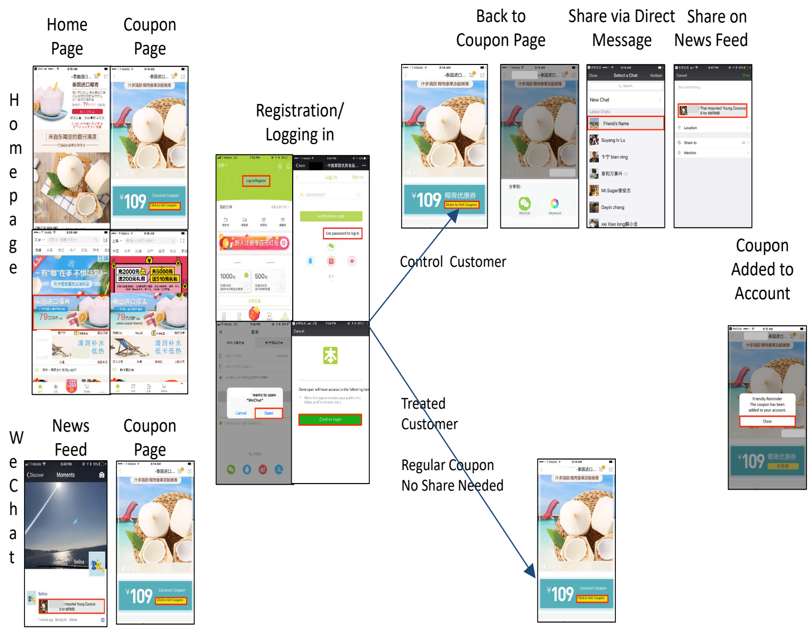

Figure A.2: Treated City Coupon Activation Steps

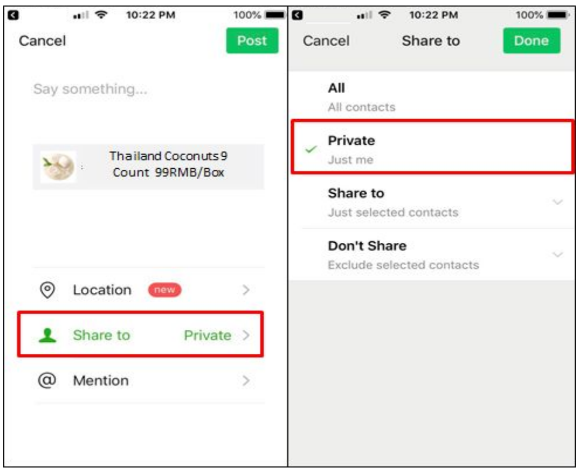

Figure A.3: Private Mode Sharing

Table A.1: Covariates comparison between all customers from 50% vs all from 0% NS cities

	Both	50% NS Cities	0% NS Cities	t.p	u.p	ks.p
Past Purchase Behavior	2					
Average Monetary Value Per Order	92.076	87.167	98.422	0	0	0
Purchased Promoted Products Before	0.343	0.332	0.356	0.009	0.009	0.088
Experiment Specific	1					
Days Since Registration	431.088	415.798	449.793	0.00000	0.013	0.0001
Response to Past Marketing Activities	1					
Average Monetary Value of Gift Cards Per Order	12.439	10.939	14.275	0.0002	0.002	0.061
Number of observations	10,830	5,959	4,871			

Notes: This table presents the set of observable customer characteristics that differ significantly at the 10% level for all three tests between all customers in 50% NS cities versus all in 0% NS cities. Columns (1), (2), and (3) report the means of individual characteristics on all customers, customers in 50% NS cities, and customers in 0% NS cities, respectively. Columns (4), (5), and (6) report the p-values of t-test, Mann-Whitney U-test, and Kolmogorov-Smirnov test for each covariate, respectively.

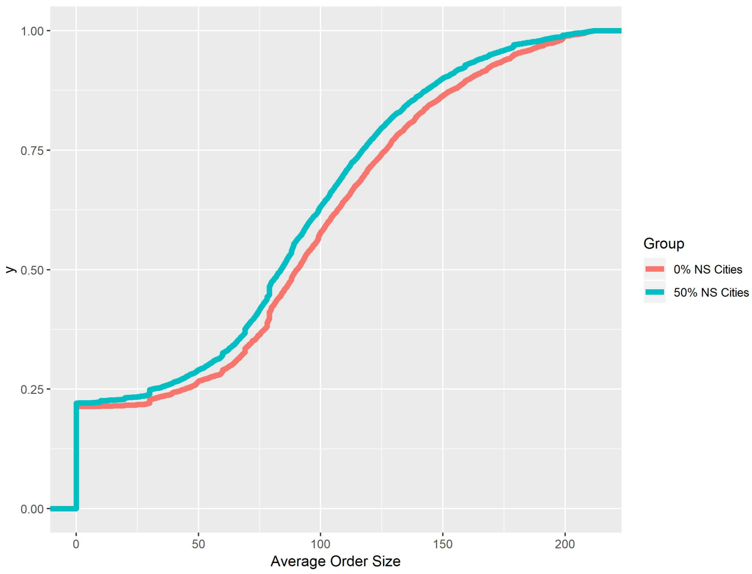

Figure A.4: Empirical cumulative distributions of customers' average order size
Notes: Among customers considering the promoted products, the average order size in the year preceding the experiment is consistently higher in 0% NS cities than in 50% cities.

A.3 Appendix Figures

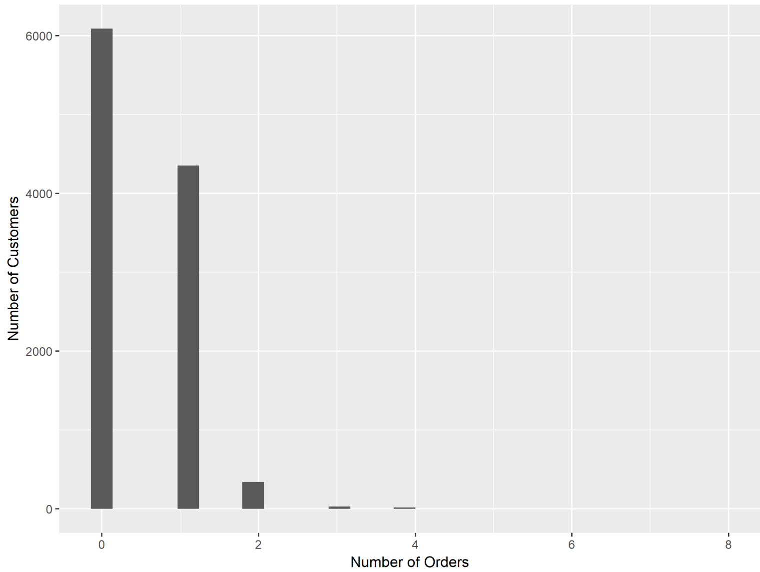

Figure A.5: Number of Orders per Customer

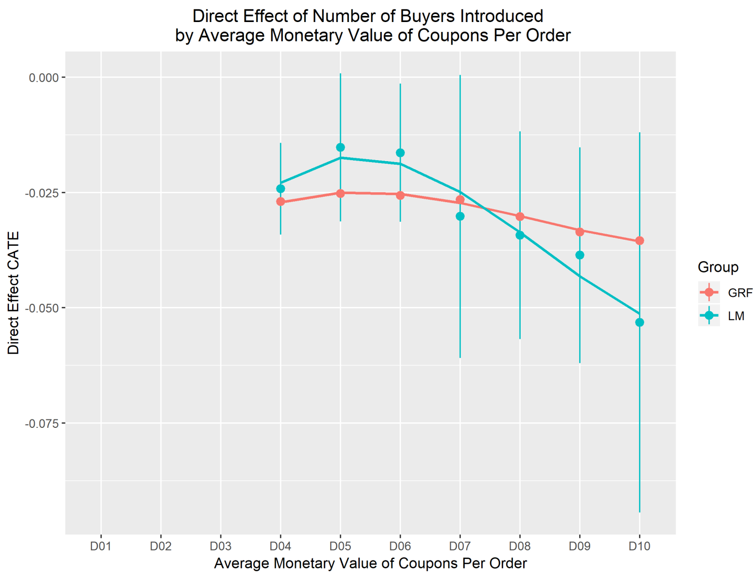

Figure A.6: Average Monetary Value of Coupons per Order

Note: This graph shows the CATEs among different segments of customers by deciles of average monetary value of coupons per order. CATEs estimated by GRF and linear model are labelled with different color. The bars are 95% confidence intervals of the CATEs. The first 3 deciles have missing values since their average monetary value of coupons per order is 0 in the past year.

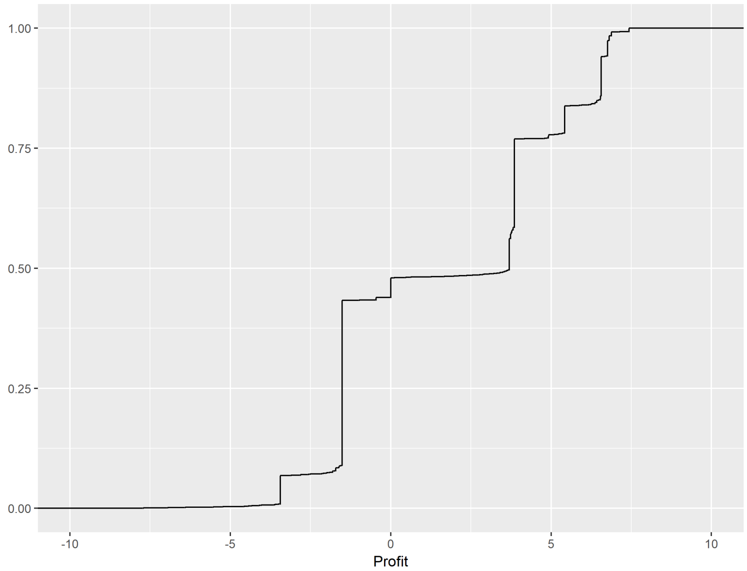

Figure A.7: Empirical cumulative distributions of promoted products' profitability
Note: Due to different wholesale costs in different cities, promoted products' profitability exhibits substantial variation, ranging from 0.6 to 1 USD (or -4 to 7 RMB).

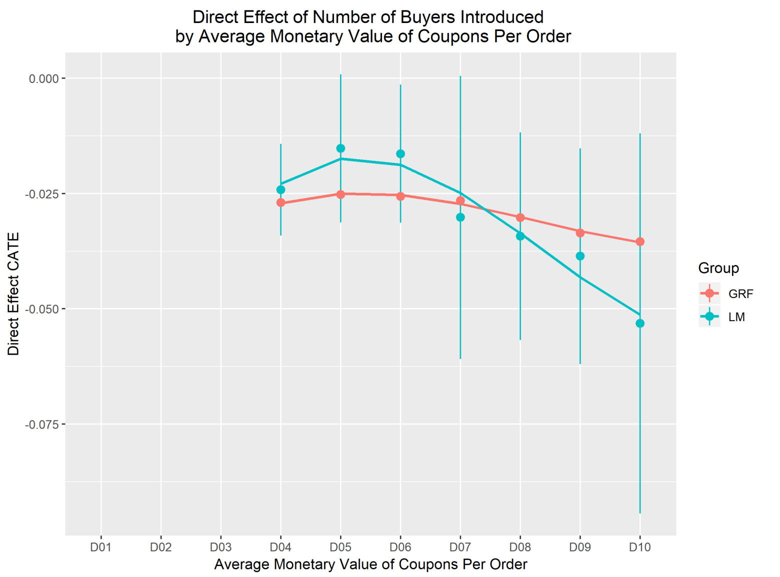

Figure A.8: Average Monetary Value of Coupons per Order
Note: This graph shows the CATEs among different segments of customers by deciles of average monetary value of coupons per order. CATEs estimated by GRF and linear model are labelled with different color. The bars are 95% confidence intervals of the CATEs. The first 3 deciles have missing values since their average monetary value of coupons per order is 0 in the past year.

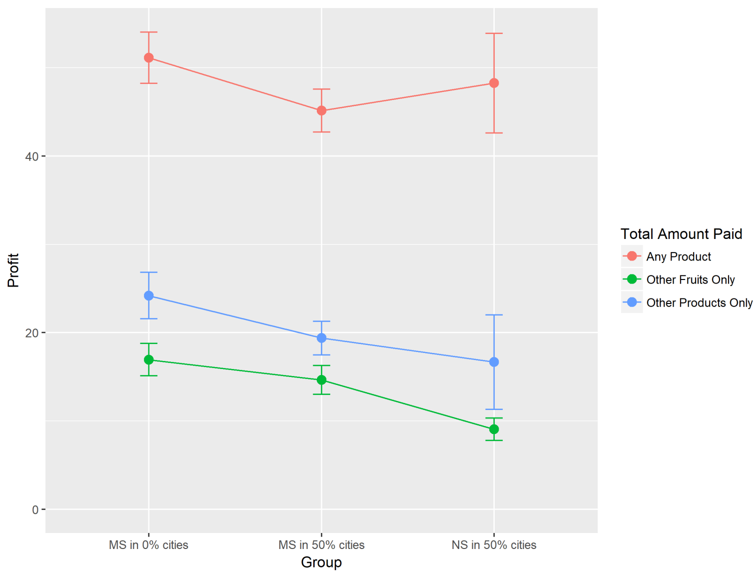

Figure A.9: Expenditure within each treatment condition

A.4 Appendix Tables

Table A.2: Heterogeneity on number of customers introduced and propensity to share (full specification)

		Number of Customers Introduced			Propensity to Share		
		CATE Greater than Median if Continuous			CATE Greater than Median if Continuous		
		Yes	No	p-value of test (1) = (2)	Yes	No	p-value of test (4) = (5)
Covariate	Median	(1)	(2)	(3)	(4)	(5)	(6)
First-Tier Cities		-0.093***	-0.057***	0.031	-0.522***	-0.460***	0.002
		(0.011)	(0.009)		(0.011)	(0.016)	
Days Since Registration	390	-0.101***	-0.061***	0.011	-0.540***	-0.462***	0.000
		(0.013)	(0.009)		(0.013)	(0.013)	
Average Monetary Value of Coupons per Order	27	-0.097***	-0.065***	0.047	-0.538***	-0.463***	0.000
		(0.013)	(0.010)		(0.013)	(0.013)	
Average Monetary Value of Coupons and Direct Discount per Order	38	—	—	—	-0.526***	-0.475***	0.006
		(—)	(—)		(0.013)	(0.013)	
Average Monetary Value per Order	99	-0.060***	-0.098***	0.031	—	—	—
		(0.010)	(0.014)		(—)	(—)	
Number of Orders with Direct Discount Last Year	1	—	—	—	-0.479***	-0.519***	0.034
		—	—		(0.013)	(0.013)	
Purchased Promoted Products Before		-0.101***	-0.071***	0.047	-0.579***	-0.463***	0.000
		(0.015)	(0.009)		(0.016)	(0.011)	
Average Monetary Value of Gift Cards per Order	0	-0.034***	-0.089***	0.019	-0.450***	-0.509***	0.027
		(0.010)	(0.009)		(0.025)	(0.010)	
Days Since Recent Purchase	20	-0.105***	-0.052***	0.002	-0.563***	-0.416***	0.000
		(0.015)	(0.009)		(0.015)	(0.014)	
Number of Orders Last Year	10	-0.061***	-0.095***	0.052	-0.456***	-0.519***	0.002
		(0.009)	(0.014)		(0.014)	(0.015)	
Active Customers		—	—	—	-0.490***	-0.542***	0.022
		—	—		(0.010)	(0.021)	

Notes: This table presents the heterogeneity results for the outcomes number of customers introduced and propensity to share. For each covariate, columns (1) and (2) report conditional average treatment effects (CATEs) estimated from the full-specification including additional lasso-selected controls (i.e. comparisons of means for S and NS customers in 50% NS cities) on the outcome number of customers introduced. For example, when the covariate of interest is recency, in columns (1) and (2), I estimate separate CATEs for those whose number of days since recent purchase is above and below the median, respectively. Column (3) shows the p-value from a test of equal conditional average treatment effects for the two segments of customers, corresponding to that of β_3. Similarly, columns (4), (5), and (6) report the results on the propensity to share outcome.

Variable Name	Importance
Days Since Registration	0.231
Days Since Recent Purchase	0.153
Number of Orders Last Year	0.145
Average Monetary Value of Coupons per Order	0.118
Average Monetary Value of Coupons and Direct Discount per Order	0.099
Average Monetary Value per Order	0.071
Average Monetary Value of Direct Discount per Order	0.047
PCScore	0.031
Purchased Promoted Products Before	0.023
Purchased Exact Promoted Bundle Before	0.021

Table A.3: **GRF Variable Importance Ranking for Number of Customers Introduced**

Variable Name	Importance
Days Since Recent Purchase	0.241
Days Since Registration	0.155
Number of Orders Last Year	0.124
Purchased Exact Promoted Bundle Before	0.113
Average Monetary Value of Coupons per Order	0.085
Purchased Promoted Products Before	0.070
Average Monetary Value per Order	0.060
Average Monetary Value of Direct Discount per Order	0.045
Average Monetary Value of Coupons and Direct Discount per Order	0.035
PCScore	0.019

Table A.4: **GRF Variable Importance Ranking for Probability of Successfully Sharing**

Variable Name	Importance
Days Since Registration	0.194
Average Monetary Value of Coupons per Order	0.150
Average Monetary Value of Coupons and Direct Discount per Order	0.113
Days Since Recent Purchase	0.090
Number of Orders Last Year	0.083
Average Monetary Value per Order	0.078
Number of Orders Using Coupons Last Year	0.053
Average Monetary Value of Direct Discount per Order	0.048
PCScore	0.043
Number of Orders with Direct Discount Last Year	0.042

Table A.5: **GRF Variable Importance Ranking for Number of Existing Customers Introduced**

Variable Name	Importance
Number of Orders Last Year	0.221
Days Since Recent Purchase	0.155
Days Since Registration	0.122
Average Monetary Value of Coupons per Order	0.083
Average Monetary Value per Order	0.081
Average Monetary Value of Coupons and Direct Discount per Order	0.076
Average Monetary Value of Direct Discount per Order	0.071
Number of Orders with Direct Discount Last Year	0.065
Number of Orders Using Coupons Last Year	0.042
PCScore	0.024

Table A.6: **GRF Variable Importance Ranking for Number of New Customers Introduced**

Variable Name	Importance
Days Since Recent Purchase	0.155
Number of Orders Last Year	0.145
Average Monetary Value of Coupons and Direct Discount per Order	0.142
Average Monetary Value of Coupons per Order	0.139
Days Since Registration	0.134
Average Monetary Value per Order	0.083
Average Monetary Value of Direct Discount per Order	0.052
Number of Orders with Direct Discount Last Year	0.043
PCScore	0.032
Number of Orders Using Coupons Last Year	0.032

Table A.7: **GRF Variable Importance Ranking for Number of Buyers Introduced**

A.5 Balance Test

Table A.8: City-level balance test

	All	Treatment	Control	t.p	u.p	ks.p
City Characteristics						
Household Registered Population of Districts (10,000 persons)	254.437	261.696	246.951	0.825	0.461	0.646
Annual Average Population of City (10,000 persons)	571.300	598.022	543.744	0.497	0.485	0.663
Annual Average Population of Districts (10,000 persons)	251.320	258.864	243.540	0.818	0.469	0.646
Natural Growth Rate of City	6.288	6.061	6.522	0.722	0.875	0.907
Natural Growth Rate of Districts	6.471	6.800	6.131	0.599	0.738	0.870
Total Land Area of Administrative Region of Districts (sq.km)	2,566.231	2,490.242	2,644.594	0.810	0.270	0.207
Area of Built Districts (sq.km)	250.154	241.848	258.719	0.810	0.559	0.973
Population Density of City (person/sq.km)	685.351	695.672	674.706	0.798	0.984	0.905
Population Density of Districts (person/sq.km)	1,018.792	887.623	1,154.060	0.076	0.238	0.402
GRP of Districts (10,000 yuan)	35,645,467.000	35,669,478.000	35,620,706.000	0.997	0.871	0.835
Per Capita GRP of City (yuan)	70,682.910	68,185.460	73,258.410	0.554	0.519	0.733
Per Capita GRP of Districts (yuan)	77,176.010	75,390.210	79,017.620	0.716	0.604	0.758
GRP Growth Rate of City	8.064	8.168	7.957	0.632	0.995	1.000
GRP Growth Rate of Districts	8.192	8.096	8.292	0.730	0.572	0.941
Primary Industry % GRP of City	7.264	8.335	6.160	0.082	0.105	0.079
Primary Industry % GRP of Districts	3.810	4.608	2.987	0.040	0.100	0.124
Secondary Industry % GRP of City	46.305	46.943	45.647	0.563	0.840	0.809
Secondary Industry % GRP of Districts	46.729	47.855	45.567	0.371	0.544	0.913
Tertiary Industry % GRP of City	44.892	44.721	45.068	0.892	0.331	0.474
Tertiary Industry % GRP of Districts	47.923	47.537	48.321	0.775	0.430	0.458
City Public Finance Income (10,000 yuan)	5,957,269.000	5,710,073.000	6,212,190.000	0.835	0.974	0.646
City Public Finance Expenditure (10,000 yuan)	7,749,142.000	7,448,789.000	8,058,881	0.826	0.901	0.923
City Expenditure for Science and Technology (10,000 yuan)	296,180.200	263,601.800	329,776.700	0.637	0.681	0.795
City Expenditure for Education (10,000 yuan)	1,275,627.000	1,241,158.000	1,311,174.000	0.849	0.912	0.993
Districts Public Finance Income (10,000 yuan)	4,770,270.000	4,463,578.000	5,086,546.000	0.799	0.943	0.997
Districts Public Finance Expenditure (10,000 yuan)	5,748,566.000	5,346,061.000	6,163,650.000	0.774	0.779	0.788
Districts Expenditure for Science and Technology (10,000 yuan)	254,099.400	216,495.400	292,878.600	0.586	0.616	0.579
Districts Expenditure for Education (10,000 yuan)	850,870.100	801,670	901,607.700	0.793	0.974	0.944
Deposits of National Banking System of City (10,000 yuan)	110,081,268.000	102,644,557.000	117,750,375.000	0.773	0.634	0.401
Deposits of National Banking System of Districts (10,000 yuan)	91,605,997.000	82,847,174.000	100,638,533.000	0.738	0.394	0.707
Household Saving Deposits of City (10,000 yuan)	36,784,212.000	36,672,513.000	36,899,403.000	0.983	0.820	0.823
Household Saving Deposits of Districts (10,000 yuan)	25,659,036.000	25,208,080.000	26,124,085.000	0.935	0.408	0.481
Loans of National Banking System of City (10,000 yuan)	69,363,661.000	67,795,862.000	70,980,453.000	0.903	0.562	0.409
Loans of National Banking System of Districts (10,000 yuan)	55,950,691.000	53,333,796.000	58,649,363.000	0.840	0.379	0.726
Number of observations	65	33	32			

Notes: Columns (1), (2), and (3) report the means of city characteristics in both treatment and control, the treatment, and control subsamples, respectively. Columns (4), (5), and (6) report the p-values of t-test, Mann-Whitney U-test, and Kolmogorov-Smirnov test on whether treated and control cities are significantly different in terms of the mean, median, and empirical distributions of each covariate, respectively.

Appendix B

Appendix for Chapter 2

B.1 Appendix on Optimal Targeting Results

Figure B.1: Comparison of Probability Density Functions

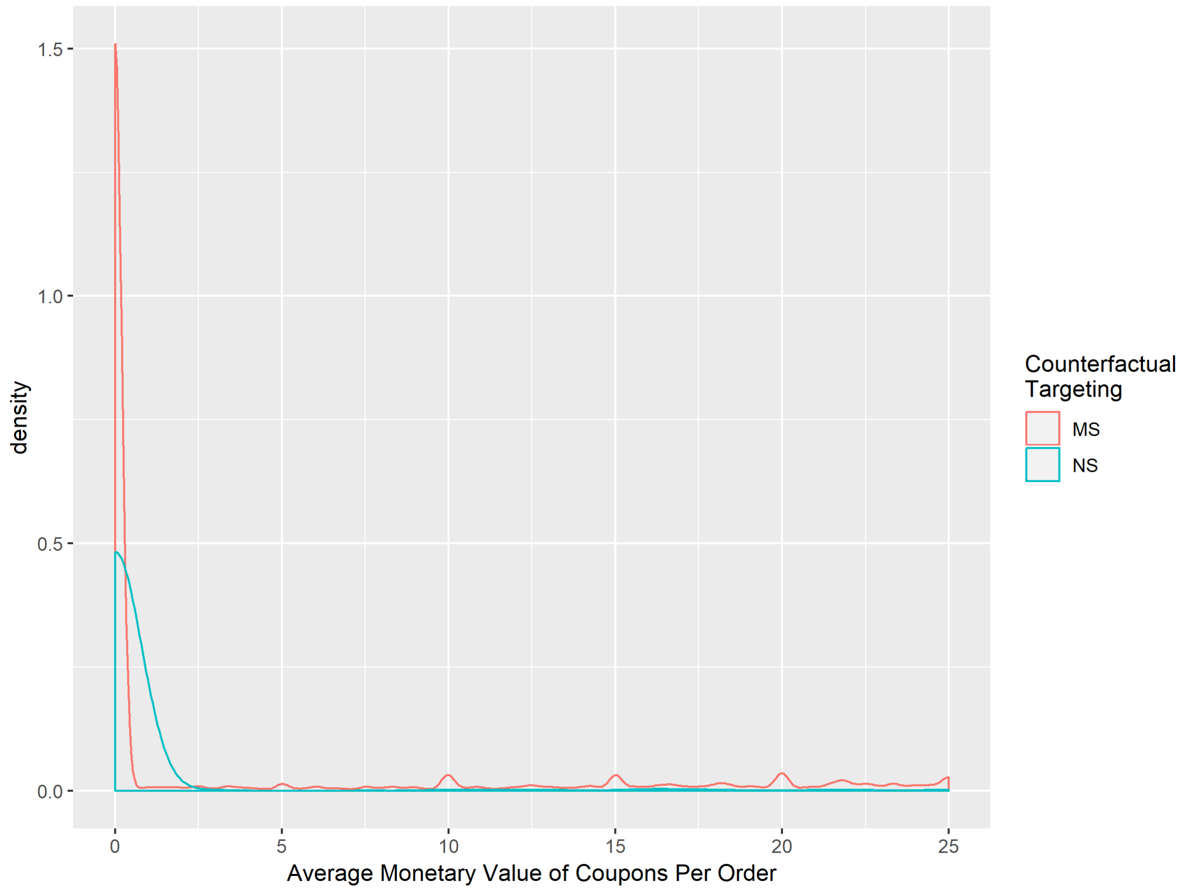

Figure B.2: Comparison of Probability Density Functions

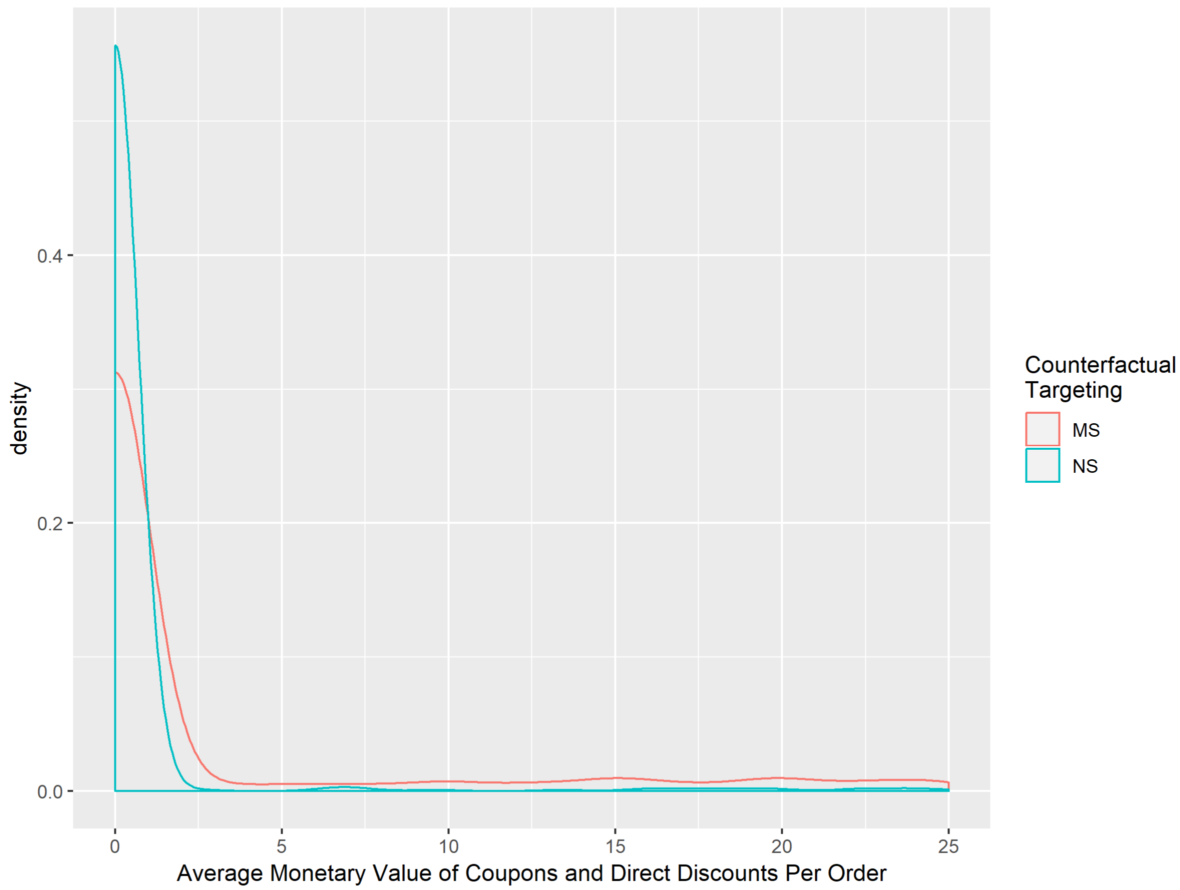

Figure B.3: Comparison of Cumulative Distribution Function

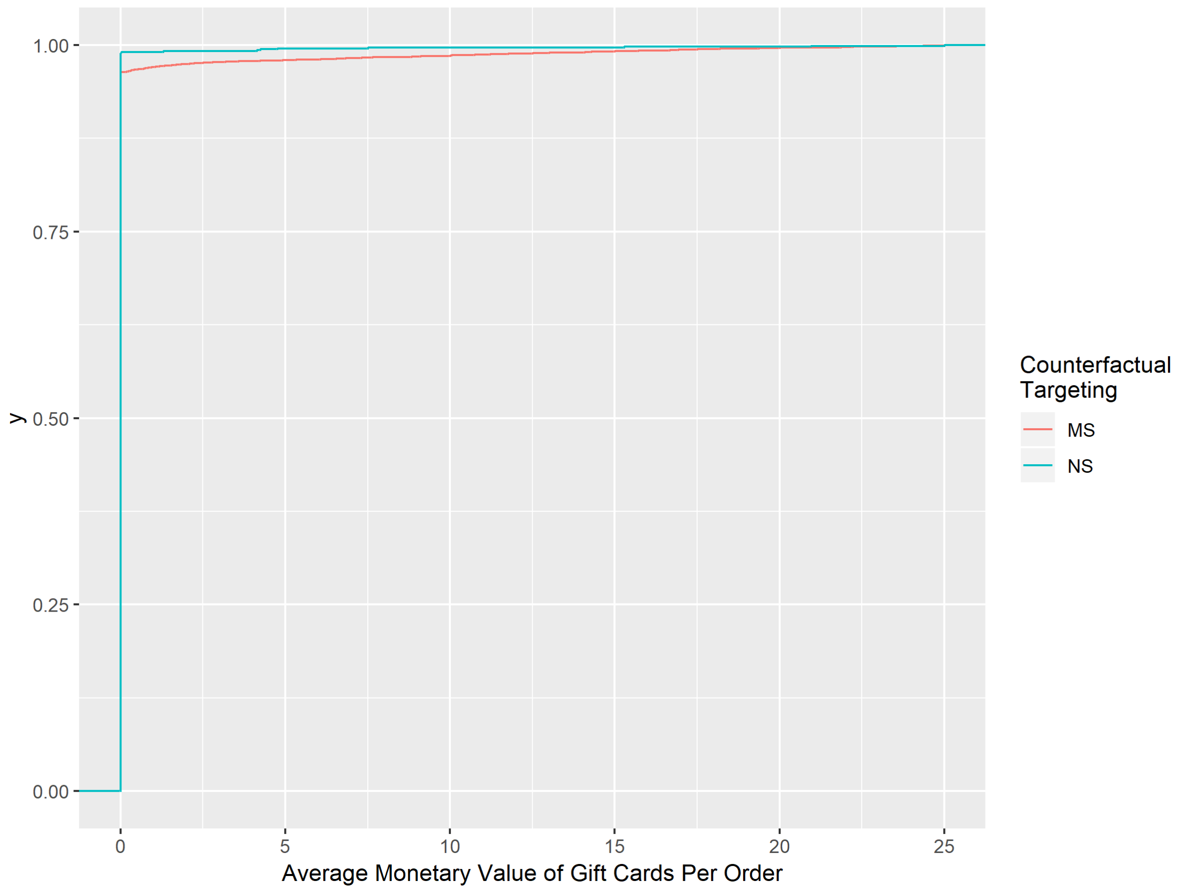

Note: This graph shows the cumulative distribution function for average monetary value of gift cards per order. Since these values are zero in most observations, cumulative distribution function can present the comparison better than probability density function.

Figure B.4: Comparison of Probability Density Functions

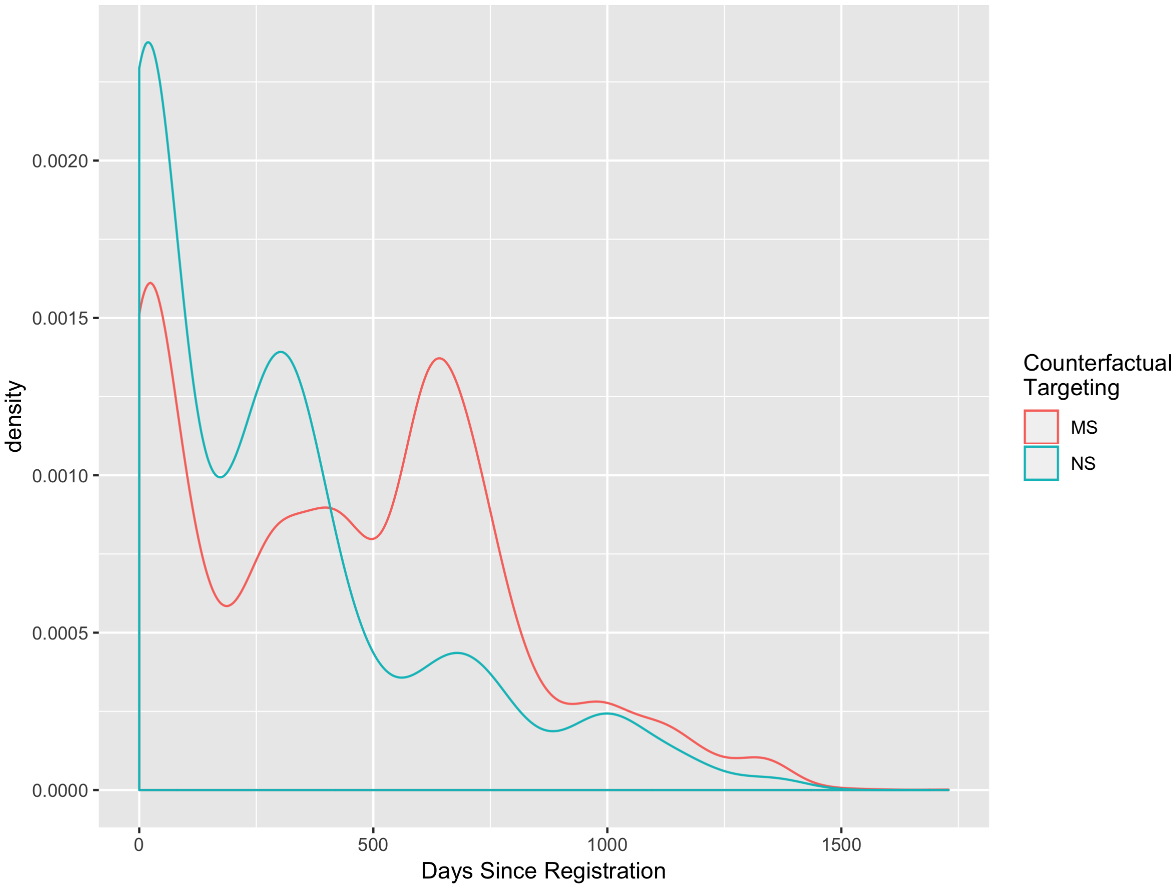

Figure B.5: Comparison of Probability Density Functions

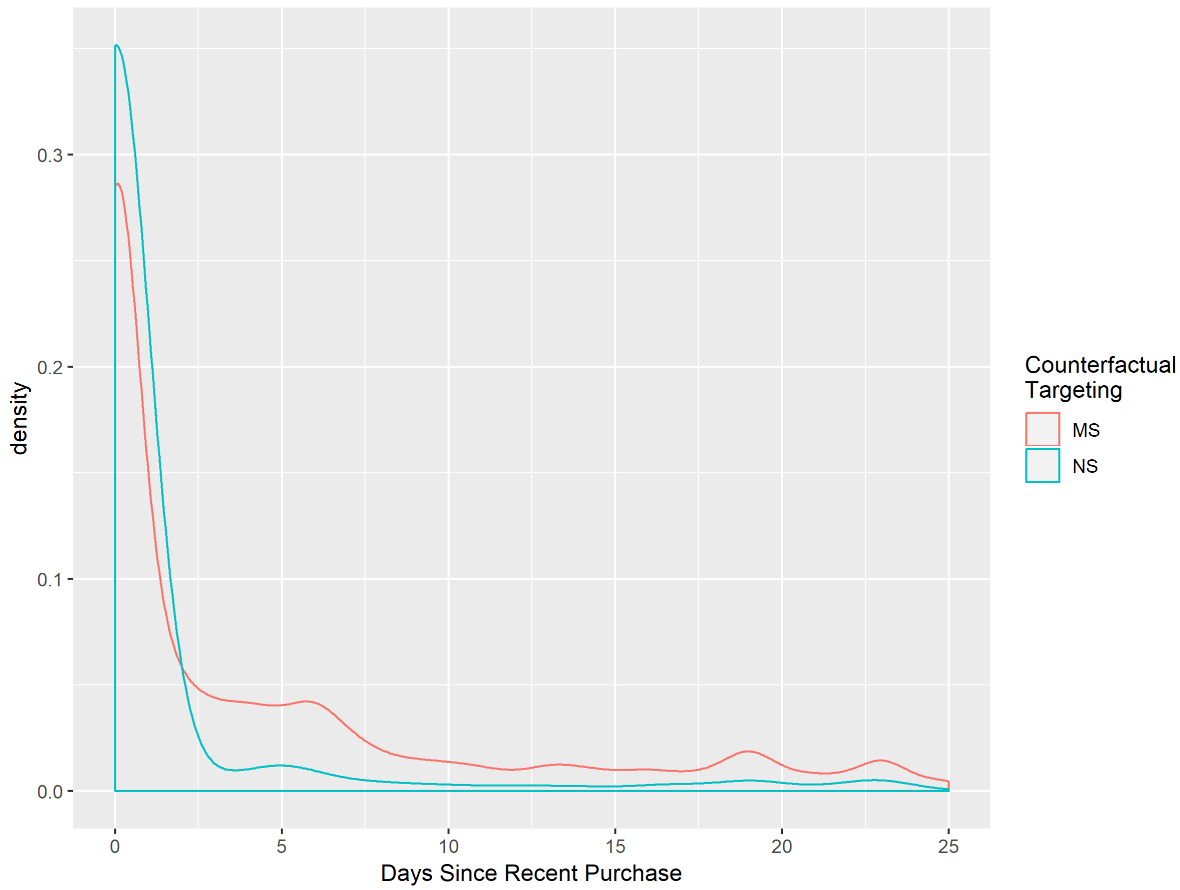

Figure B.6: Comparison of Probability Density Functions

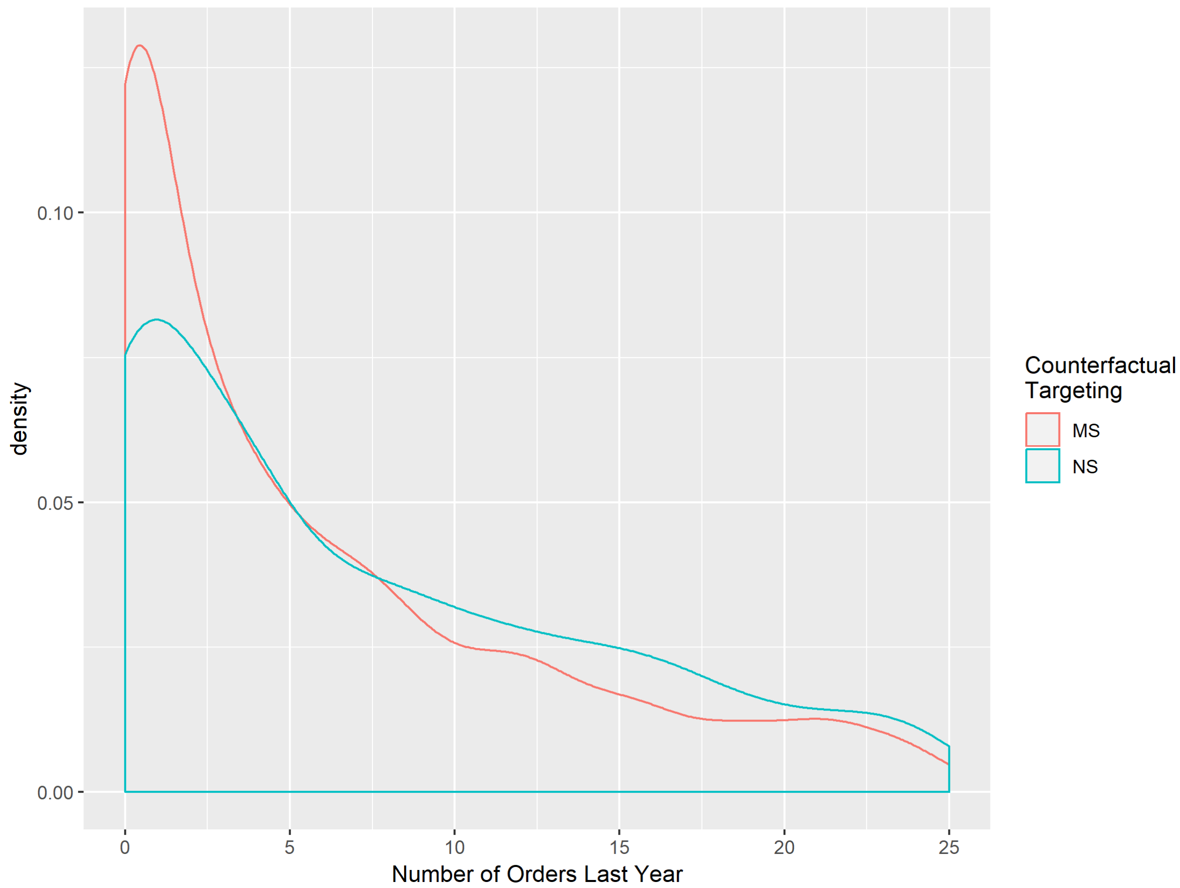

Figure B.7: Comparison of Probability Density Functions

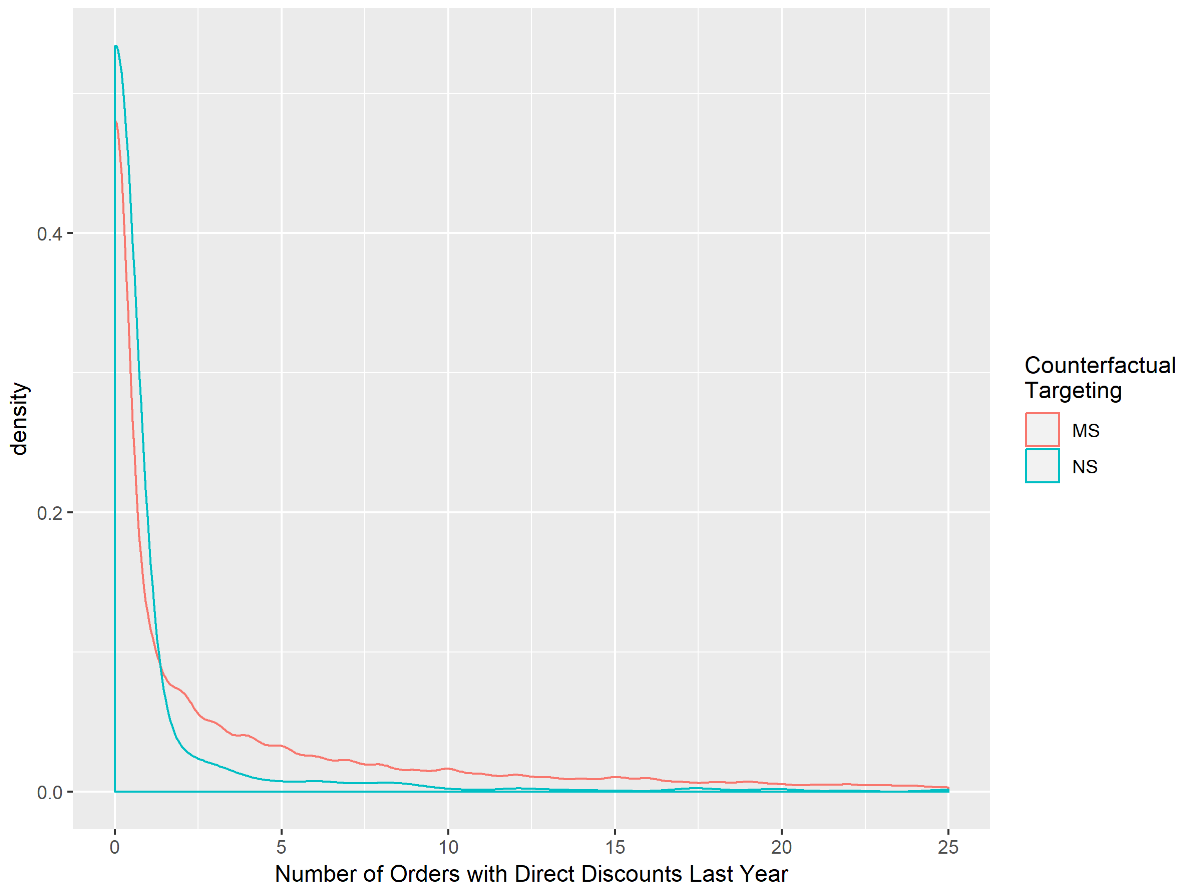

Figure B.8: Comparison of Probability Density Functions

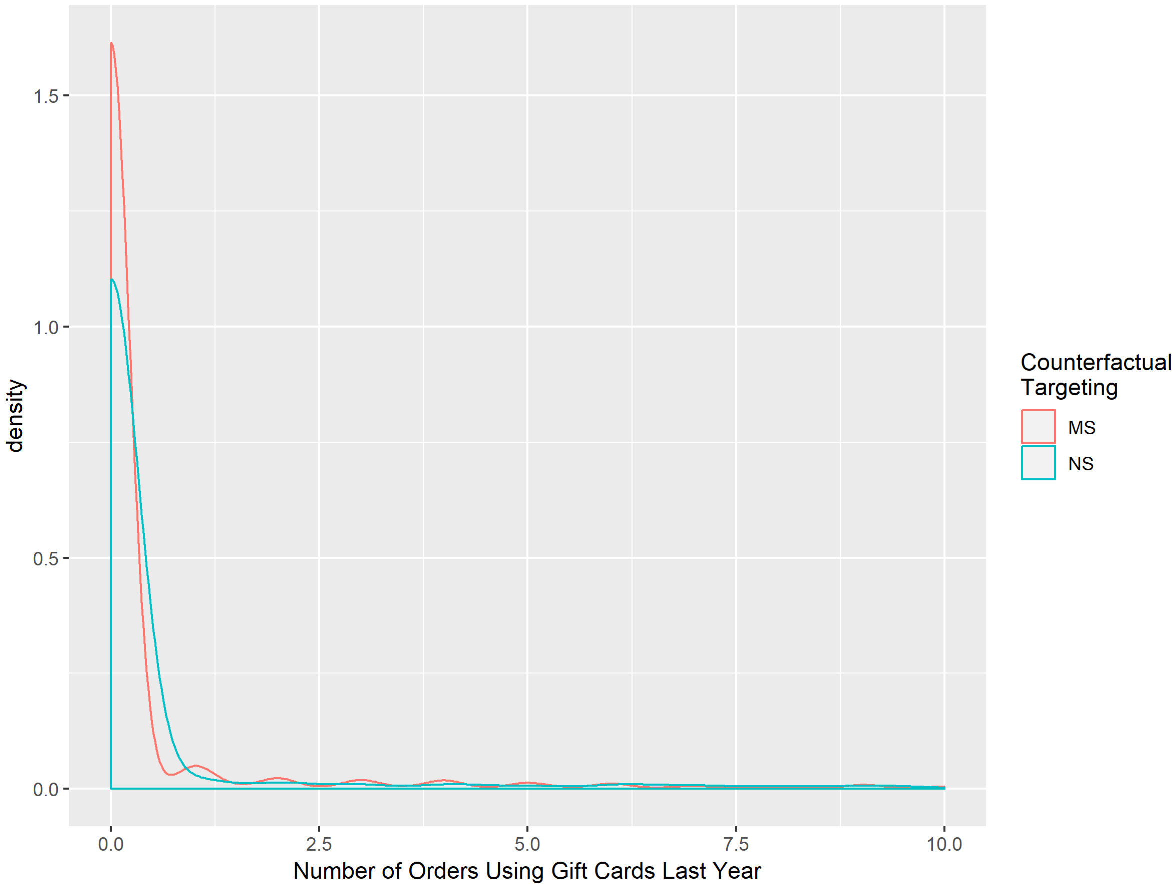

Figure B.9: Counterfactual Assignment under Optimal Targeting

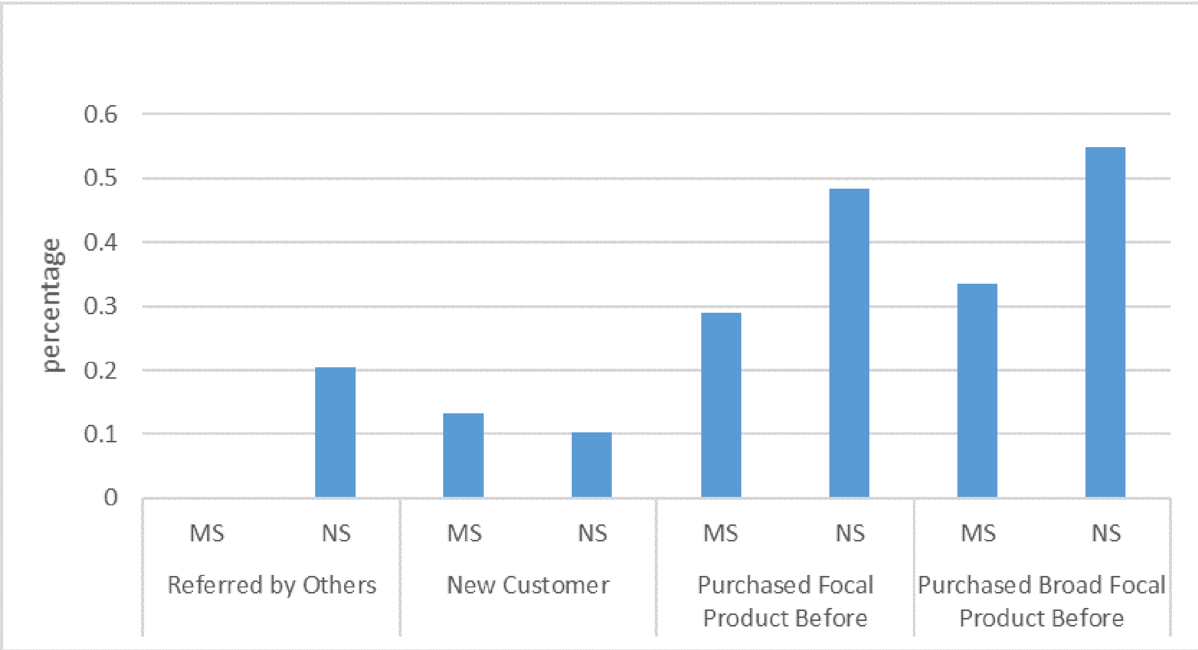

B.2 Appendix on Profitability

Figure B.10: Estimated average adjusted profit

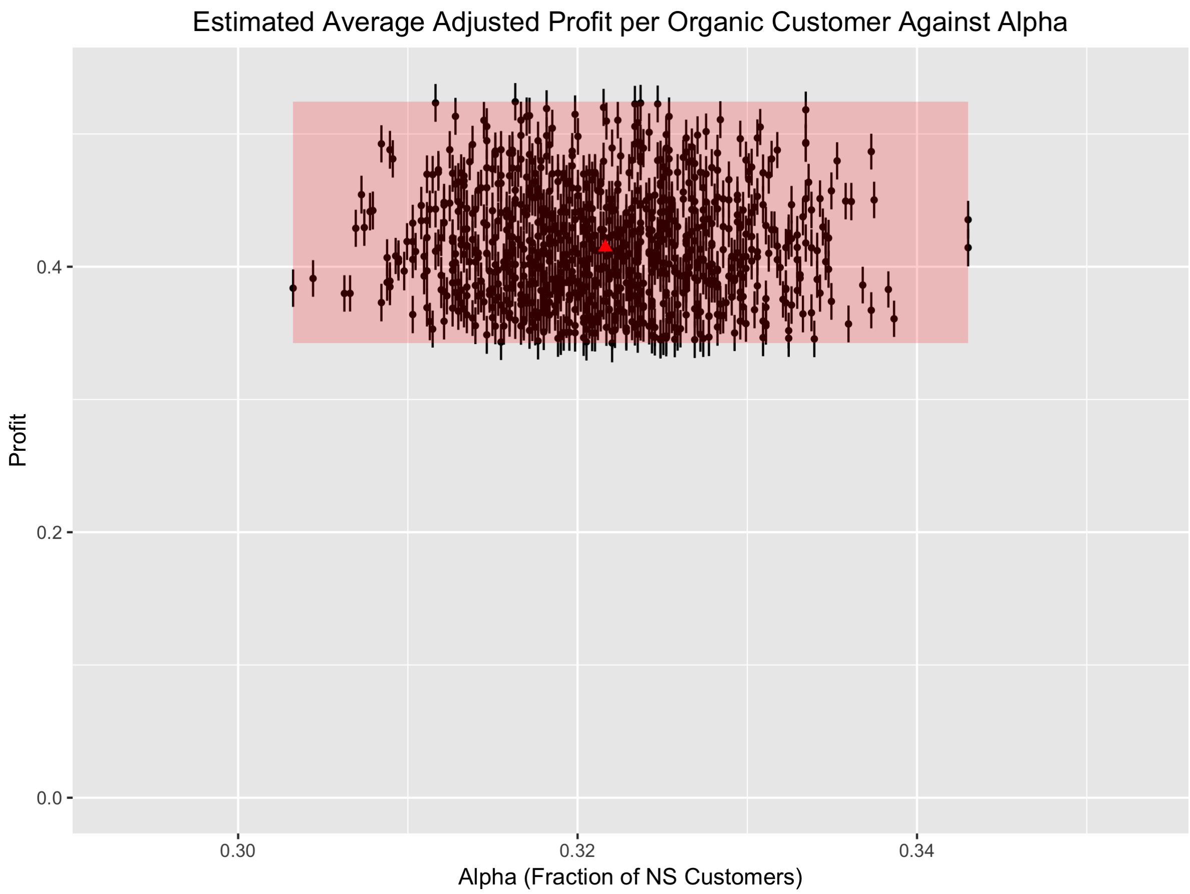

Figure B.11: Estimated percent of increase of average adjusted profit lift factor

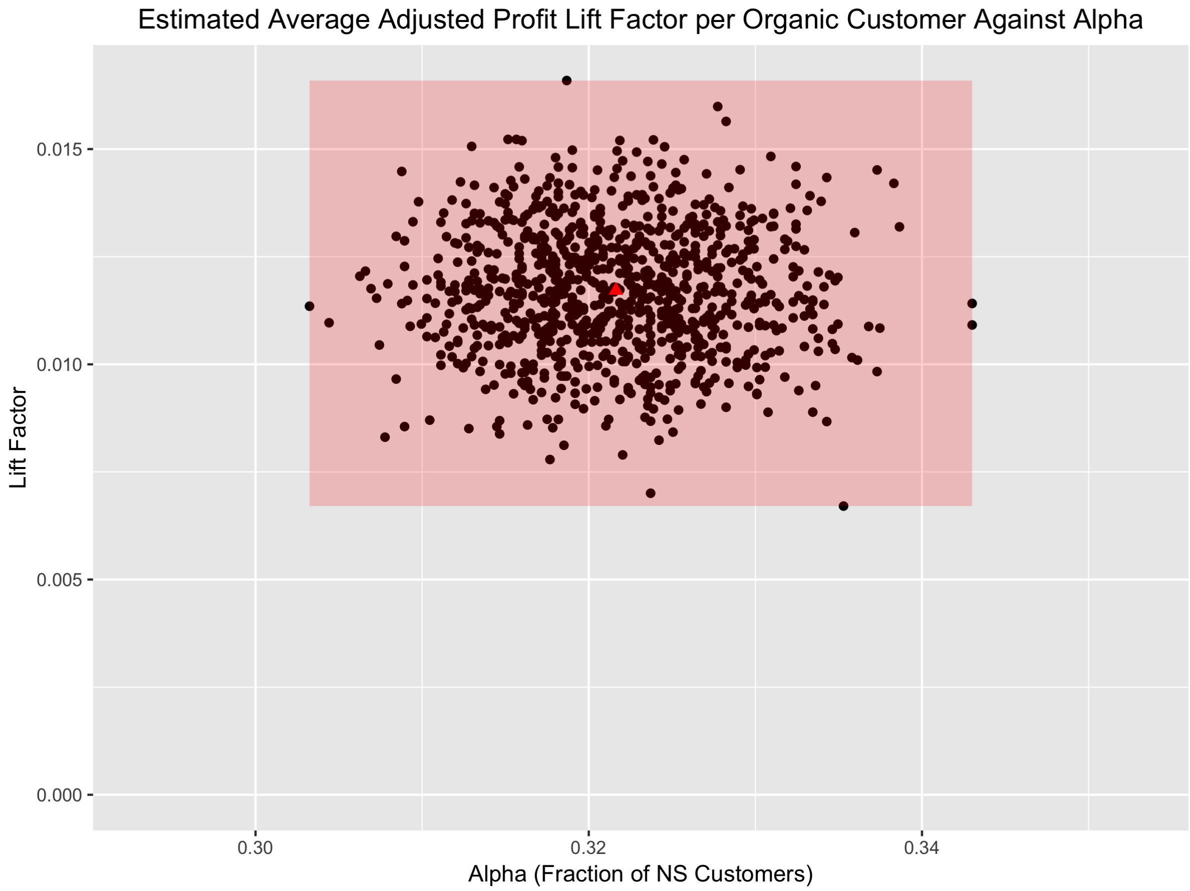

Figure B.12: Estimated adjusted profit

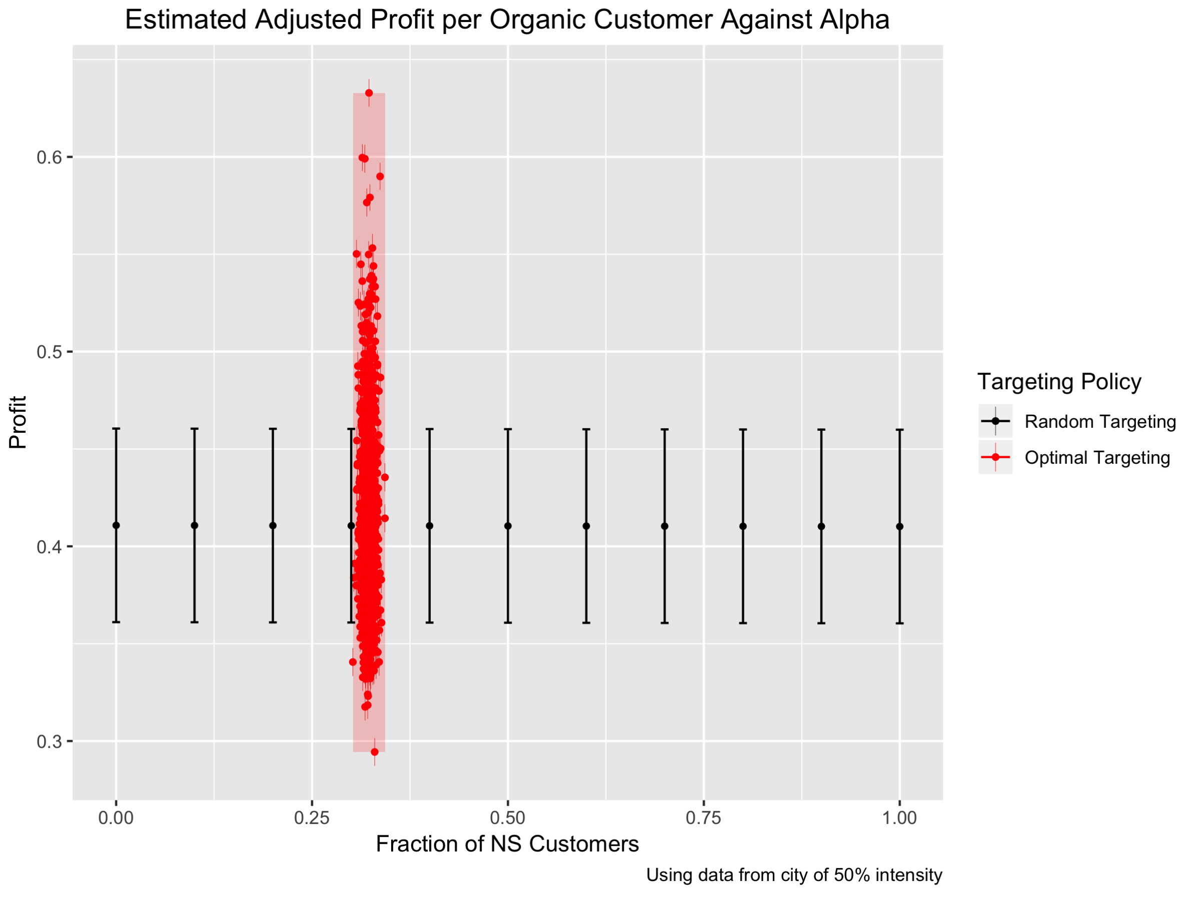

Figure B.13: Estimated adjusted profit lift factor