

THE UNIVERSITY OF CHICAGO

INVESTIGATING THE EFFECTS OF INCLUDING DISCOUNT INFORMATION IN
ADVERTISING

A DISSERTATION SUBMITTED TO
THE FACULTY OF THE UNIVERSITY OF CHICAGO
BOOTH SCHOOL OF BUSINESS
IN CANDIDACY FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

BY
SHIRSHO BISWAS

CHICAGO, ILLINOIS

JUNE 2020

Copyright © 2020 by Shirsho Biswas

All Rights Reserved

This dissertation is dedicated to the memory of my beloved mother and closest friend, Antara Biswas. Her love and care have been instrumental in shaping me into the person I am and making this document possible. It is also dedicated to my father, Siddhartha Biswas, and my brother, Shaurjo Biswas, who got me started on the journey to my Ph.D. and encouraged me through it. I love you with all my heart.

TABLE OF CONTENTS

LIST OF FIGURES	v
LIST OF TABLES	vi
ACKNOWLEDGMENTS	vii
ABSTRACT	viii
1 INTRODUCTION	1
2 EMPIRICAL CONTEXT AND EXPERIMENT DESIGN	10
2.1 Data	18
3 EFFECT ON DEMAND	20
4 EFFECT ON SEARCH	25
4.1 Self-selection of consumers into search	30
5 EFFECT ON CONVERSION TO PURCHASE CONDITIONAL ON SEARCH	34
5.1 Matching on observables using a causal forest	35
5.2 Modeling the consumer decision process to account for selection on unobservables and observables	38
5.2.1 Model specification	39
5.2.2 Distributional Assumptions	43
5.2.3 Estimation Methodology	44
5.2.4 Results	45
5.3 Ruling out alternate explanations	46
6 POST-EXPERIMENT DEMAND	48
7 OPTIMAL TARGETING	51
8 MANAGERIAL IMPLICATIONS	53
9 CONCLUSION	55
REFERENCES	57
APPENDICES	59
A RANDOMIZATION CHECKS	59

LIST OF FIGURES

1.1	Papa Murphy [®] Discount Advertisement	1
1.2	Illustration of research question	2
1.3	Effects of different ways of communicating brand and discount information driving search and purchase decisions	4
2.1	Different Ad Creatives	12
2.2	Screenshot of ad push notification on an Android phone and the landing page after clicking on the ad	13
2.3	Experiment Design	14
2.4	Steps in the purchase funnel	15
2.5	Listing and Search pages	16
2.6	Menu and Cart pages	16
3.1	Effect of the different types of ads on demand for the focal restaurant	20
4.1	Effect of the different types of ads on probability of visiting the focal restaurant menu page	27
4.2	Effect of different types of ads on the probability of visiting the focal restaurant menu page on sub-samples of consumers split by expectations of discount conditional on receiving a notification	28
4.3	Effect of the different types of ads on probability of visiting the focal restaurant cart page	29
4.4	Pre experiment customer ‘characteristics’ for individuals who chose to visit the focal restaurant menu page under different discount and ad conditions (1)	31
4.5	Pre experiment customer ‘characteristics’ for individuals who chose to visit the focal restaurant menu page under different discount and ad conditions(2)	32
5.1	Distribution of differences in probability of conversion at cart after having been treated with a price ad and brand ad	38
5.2	Distribution of differences in probability of conversion at cart after having been treated with a price ad and brand ad for different discount levels	39
5.3	Distribution of differences in probability of conversion at cart after having been treated with a price ad and brand ad for previous, new and experienced focal restaurant customers	40
6.1	Post experiment demand from focal restaurant	49
6.2	Post experiment demand from focal restaurant for consumers who did not search or purchase from it during expt.	50
7.1	Revenue Curves with different advertising types	51

LIST OF TABLES

2.1	Difference between total number of notifications randomly assigned and actually received	19
3.1	Effect of Ads on Demand for the Focal Restaurant	23
3.2	Meal value conditional on preparing a cart for the focal restaurant	24
5.1	Three stage trivariate probit model estimation results	45
5.2	Three stage trivariate probit model estimation results with reference term added in the last stage	46
A.1	Randomization Checks 1	59
A.2	Randomization Checks 2	59

ACKNOWLEDGMENTS

I am extremely grateful to my advisor, Pradeep Chintagunta, and the members of my dissertation committee, Sanjay Dhar, Jean-Pierre Dube and Anita Rao for their constant support and guidance. I am very thankful to Pradeep and Sanjay for their friendship and mentorship right from the start of my PhD journey to its end, for teaching me, and for investing so much of their time listening to my many different ideas and questions and providing feedback and answers. I thank JP and Anita for their thoughtful and constructive inputs on this project and others. I also thank the other members of the Marketing Faculty at Booth for their advice and help in preparing me for the job market. I thank my friends and colleagues from the PhD program at Booth, who have helped me through the program and made my journey more enjoyable and fun.

I thank the members of the Marketing and Data Science groups at an anonymous food delivery company for helping me implement the experiment and collect data for my dissertation.

Finally, I would like to thank my family - my father and brother for their constant love, support, advice and encouragement, which kept me going forward; my sister-in-law, Debaleena Basu, for raising my spirits and making me believe in myself during hard times; and my niece, Sia, for always making me smile, through good times and bad.

ABSTRACT

Consider a display ad for a new pizza brand that announces a discount. Such discount advertising mentions the brand and often shows a picture of the product, thus informing consumers of the brand and some characteristics (the ‘brand advertising effect’). It also highlights a discount, thus informing consumers of the existence of a discount (the ‘pure promotion effect’). There could be a third effect, the ‘discount spotlighting effect’ - a reduction in brand preference from discount advertising that arises from the brand *choosing* to highlight a discount in its advertising and marketing itself based on low price. These three effects are typically confounded and, in particular, the ‘discount spotlighting’ effect has not been studied before. Two identical consumers, one who has seen a brand ad and then found out about the discount separately, and another one who has seen a discount ad, have the same information about the brand and the discount. However, they will have different probabilities of purchase due to the ‘discount spotlighting’ effect. I demonstrate the existence of this effect by designing and implementing a field experiment on a food delivery app with exogenous variation in advertising intensity, the presence of discount information in ads, and discount level for a focal restaurant. More broadly, I use this experiment to investigate the differences in the effects of ads with and without explicit discount information (‘discount ads’ and ‘brand ads’) on demand for the advertised brand and the different stages of the purchase funnel (search and purchase). I show that discount ads can induce more consumers to search for the product, but result in lower conversion to purchase conditional on search for a given consumer, relative to brand ads. Further, I show that targeting different consumers with brand ads or discount ads along with targeted discounts based on their probability of search and conditional purchase increases firm revenues relative to a single, optimal type of ad and discount level.

CHAPTER 1

INTRODUCTION

We see many online advertisements (like the Papa Murphy[®] display ad below in Figure 1.1) which contain brand information (like the name of the brand and a picture of the product) and *also* highlight the presence of a price discount.



Figure 1.1: Papa Murphy[®] Discount Advertisement

Previous literature, both empirical (Mela et al., 1997; Jedidi et al., 1999; Kaul and Wittink, 1995; Blattberg and Neslin, 1989; Rao and Monroe, 1989) and theoretical (Scitovszky, 1945; Olson, 1977; Dodson et al., 1978; Aaker, 2009) have discussed potential negative effects of price promotions or discount advertising in terms of reducing brand equity and making consumers more price elastic (while brand oriented advertising has the opposite effect). However, the literature studying the unintended negative effects of promotional advertising has failed to distinguish between two possible sources of these negative effects:

1. The consumer makes a negative quality inference about the product because they see that it is on promotion (Scitovszky, 1945; Olson, 1977). This effect would exist regardless of whether the consumer found out about the promotion from an advertisement or elsewhere (say from browsing through the shelf in a supermarket or an online store and noting available price discounts)
2. The consumer makes a (second or additional) negative quality inference about the product because they learn about the discount *through an ad* (as opposed to other means like browsing through the shelf in a supermarket or an online store and noting

available price discounts). Advertising is a medium by which firms highlight their differentiating features and the strategic decision of the firm to build an ad campaign around a discount may lead to an additional negative quality inference on top of source 1 above. This effect only exists if the consumer learns about the discount through an ad.

While the first source above has been theoretically proposed (Scitovszky, 1945; Olson, 1977) and empirically confirmed (Blattberg and Neslin, 1989; Rao and Monroe, 1989) in literature that studies price promotions (and not promotional advertising), the second source has not been raised in the literature. Papers that have studied the effects of promotional advertising (Mela et al., 1997; Kaul and Wittink, 1995) have conflated the two possible sources. This second potential source of the negative effects of discount advertising forms the basis of the research question of this paper (pictorially represented in Figure 1.2): Does additional brand dilution ¹ occur, when consumers learn about a discount *through an ad* vs. if see an ad containing only brand information and they learn about the discount through other means?



Figure 1.2: Illustration of research question

Note that the negative quality inferences discussed above may only exist and be relevant in contexts where there is some uncertainty about product quality in the consumer's mind i.e. it likely does not exist for established brands for which consumers have a good idea of

1. Brand dilution here refers to a negative effect on brand preference or brand equity through a negative quality inference

product quality or a well established brand preference ². To clarify further, consider the different effects that an a discount ad like the one in Figure 1.1 may have. Such ads may have three distinct effects:

1. The brand advertising effect - this is the effect of informing consumers of the existence of the brand (Nelson, 1970) and may also increase brand preference (Comanor and Wilson, 1979). This effect is common to both brand and price advertising, and is presumably positive for the advertised brand at all stages of the purchase funnel, i.e. awareness, search and conditional purchase.
2. The pure promotion effect - this is the same effect that would exist if the consumer saw a shelf discount of the same amount at the store. A large discount (relative to expectations) will increase search (for more product information), and will increase conditional purchase due to the lower price. However, knowing that the brand is offering a discount could also affect brand preference negatively (Scitovszky, 1945; Olson, 1977; Blattberg and Neslin, 1989) and this might lower search and conditional purchase.
3. The ‘discount spotlighting’ effect - a consumer may infer that a firm that chooses to highlight a discount in its advertising is not differentiated in any way apart its low price, and hence is marketing itself based on price. This would negatively affect brand preference and would lower search and conditional purchase. This is an effect that price advertising has, but not brand advertising. Also, this effect is different from the pure promotion effect as it arises from the fact that a discount is *highlighted in an ad*, and not from the fact that the discount exists. It is unique to discount advertising and does not occur when the same discount information is learned by consumer, say, at the store.

2. I show evidence of this in my empirical context below

In terms of the two possible sources of unintended negative consequences of discount advertising mentioned above, the first one is under the ‘pure promotion effect’ and the second one is the ‘discount spotlighting effect’. Thus, to rephrase the main objective of the paper, we want to test whether the ‘discount spotlighting effect’ exists or not. How these three different effects influence a consumer who either (1) sees a brand ad and then learns of a discount at the store or (2) sees a discount ad with both brand information and discount information are shown in Figure 1.3.




Figure 1.3: Effects of different ways of communicating brand and discount information driving search and purchase decisions

It is important for marketing academics to know whether the ‘discount spotlighting effect’ exists (in addition to the ‘brand advertising effect’ and the ‘pure promotion effect’) for two reasons. First, if it exists, we learn that previously documented unintended negative consequences of discount advertising (Mela et al., 1997; Jedidi et al., 1999; Kaul and Wittink, 1995; Blattberg and Neslin, 1989; Rao and Monroe, 1989) arise not only from the direct effect of informing consumers of a discount, but additionally also from doing so *through an ad*. Thus, we learn about an additional negative effect of discount advertising which hasn’t been discussed in the literature. Second, its existence would raise an important discussion about the optimal way for firms to communicate brand and discount information to the consumer : whether to (a) bundle these together into a discount ad or (b) separately communicate

them, first through a brand ad and then letting the consumer find out about the discount at the store. Consider the problem of getting the most number of people to purchase the product during the promotional campaign period (as opposed to over the long term or after the promotional period, which has been the focus of most papers studying unintended negative consequences of promotions and promotional advertising (Mela et al., 1997; Jedidi et al., 1999; Kaul and Wittink, 1995; Blattberg and Neslin, 1989)). The optimal strategy for revenue maximization during the campaign depends on the relative magnitudes and signs of the ‘brand advertising effect’, ‘pure promotion effect’ and ‘discount spotlighting effect’ on the search and conditional purchase stages of the purchase funnel for each consumer. The brand advertising effect is always positive. The net effect of the ‘pure promotion effect’ on search and purchase during the promotional campaign is presumably positive in most cases. However, the discount spotlighting effect is negative. For a given consumer, consider two different scenarios as depicted in Figure 1.3. If they are shown a brand ad without any discount information, they are likely less likely to initiate search relative to a discount ad with discount information (assuming that the sum of the pure promotion effect and discount spotlighting effect on search is positive). However, they are also more likely to convert to purchase conditional on search and seeing the same discount information (since the sign of the discount spotlighting effect on conditional purchase is negative). Thus, the total effect on probability of purchase during the promotional campaign is unclear. If there were no discount spotlighting effect, then as long as showing a discount ad leads to higher probability of initiating search, the discount ad is the optimal ad to show any consumer for a fixed discount level. However, if the discount spotlighting effect exists, the negative effect on conversion to purchase relative for discount ads relative to ‘brand ads + discount information through search’, makes it hard to theoretically specify an optimal strategy, and makes it an empirical question.

The second issue above is especially important in the online context. Unlike in the offline world where brand advertising is done through eg. TV and discount advertising is done

through eg. newspaper inserts, and these are separate and serve different purposes, online ads are limited to an image to be displayed to the consumer, an email subject like or a short pop-up message on the phone. In the online context, brand ads and discount ads are substitutes for each other. Whether the firm should use the limited advertising space to highlight a discount is not clear. Firms invest a lot in optimizing their advertising content (Bertrand et al., 2010; Sudhir et al., 2016), but the dimension of whether to include discount information or not, has surprisingly been ignored; the assumption presumably being that more information about a discount is better. Notwithstanding ‘conventional wisdom’ of firms such as Google that highlighting discount information is good practice for display ads³, studying the effects of the inclusion of discount information is an important and open question for researchers and practitioners.

In this paper, I run a field experiment at an online food delivery platform where I randomly assign around 150,000 consumers to receive either a brand ad or a discount ad advertising a particular restaurant on the platform through the medium of mobile push notifications.⁴ These consumers are also randomly assigned to one of five different discount levels (0%, 10%, 20%, 30% or 40%) and the assignment of the type of ad and discount level are done independently of each other. I observe consumer decisions at both the search and purchase stages of the purchase funnel. The identification strategy is to test for the presence of the discount spotlighting effect by checking if a given consumer is less likely to purchase conditional on having searched, if they are randomly selected to receive a discount ad vs if they are randomly selected to receive a brand ad and then they see the discount after engaging in search.

First, I show through the experiment, that the probability of search is higher if a discount is mentioned in the ad (and the discount mentioned is higher than what the consumer is

3. <https://support.google.com/google-ads/answer/1722134?hl=en>

4. The brand information was kept identical for the two ads. The discount ad contains additional discount information. A third creative (intermediate ad) that only mentioned that there was a discount available without mentioning the exact percentage discount and without highlighting it was also used. The effects of this ad were very similar to the brand ad

used to seeing in these ads on the platform ⁵). Next, I show that consumers self-select into search based on the discount information provided in the ad and their expectations of discount on the platform (if no discount information is mentioned). This self-selection into search means that directly comparing conditional rates of purchase post search across ad conditions is not valid. Self-selection needs to be accounted for in order to make a fair comparison between ad conditions for a ‘given consumer’. I do this in two ways: (a) Using rich historical data on consumer search and purchase behavior on the platform and for the focal restaurant, along with an assumption of selection on *observables*, I employ a non-parametric machine learning approach in the form of a causal forest (Wager and Athey, 2018) to match on a large set of observable characteristics. (b) I allow for selection on *observables and unobservables* by specifying a multi-stage model of self selection into the various steps of the customer journey and allowing error terms across stages to be correlated. Each of these two approaches has its pros and cons as discussed in Chapter 5. Note that there is no way to experimentally avoid the self-selection problem, and the rich historical data on consumers makes this a particularly suitable context for employing recent non-parametric machine learning matching methods. Using both approaches, I find that a consumer who receives discount advertising has a lower post-search probability of purchase relative to brand advertising, thus confirming the existence of the discount spotlighting effect. Moreover, I show that in my empirical context, discount advertising in combination with a high discount, discount advertising is still the revenue maximizing strategy for the firm. Lower conversion to purchase is offset by the fact that highlighting a large discount in discount advertising attracts a large number of consumers into searching for the product. This explains the ubiquitous use of discount advertising, despite the possible negative consequences raised by this paper and previous literature. However, there may exist contexts where this is not the case, depending on the relative sizes of the ‘brand advertising effect’, ‘pure promotion effect’ and ‘discount spotlighting effect’. Managers of individual firms must experiment with

5. consumer expectations of discount in this context are discussed in more detail in Chapter 4

different ad creatives to find the optimal strategy for them. With the possibility of individual targeting, the optimal strategy may be to have both creatives, and target them at different individuals⁶.

Note that apart from the mechanism described above, there may be some alternate explanations for the finding that discount advertising results in a lower post-search probability of purchase relative to brand advertising. The first is a reference price effect. Second, consumers' price sensitivity increases. In other words, the co-efficient on price, conditional on search may be larger in absolute value under discount advertising compared to brand advertising. Third, brand information in the ad may be less salient than in the ad if accompanied by discount information. I use heterogeneity in the difference in post search conversion rates between brand ads and discount ads at different discount levels and for consumers with different levels of experience with the focal restaurant to rule out these alternate explanations⁷.

This paper makes four main contributions to the literature: (1) it contributes to the literature studying unintended negative consequences of promotional advertising (Kaul and Wittink, 1995) by demonstrating the existence of the 'discount spotlighting effect', thus documenting a source of the negative effects that hasn't been discussed before. (2) Unlike previous literature that has focused on negative consequences of promotional advertising in the long term (Mela et al., 1997; Jedidi et al., 1999), this paper documents a 'real-time', within-campaign negative effect of discount ads relative to brand ads. (3) It is the first paper to experimentally compare effects of ads with and without discount information, *holding all brand information fixed*, and thus isolating the effect of discount information in ads, on overall demand *and* on the different stages of purchase funnel. (4) It is the among the first papers to document the effects of push notification ads on sales. While this is an emerging advertising medium, there is little empirical documentation of the effects of push notification ads on demand.

6. I show that this is the case in my empirical context in Chapter 7

7. These alternate explanations are discussed in more detail in Chapter 5

The remainder of this dissertation is organized as follows. In Chapter 2, I describe the empirical setting and experimental design. In Chapter 3, I present results on the impact of the different ads on overall demand. In Chapter 4, I examine the decision to search for the advertised brand and how the different ads affect this process. In Chapter 5, I examine the conditional purchase decision post search. In Chapter 6, I look at post-experiment demand. In Chapter 7, I examine the optimal targeting strategy for the firm in terms of both advertising type and discount level. In Chapter 8, I present managerial implications and in Chapter 9, I conclude by summarizing the key findings and directions for future research.

CHAPTER 2

EMPIRICAL CONTEXT AND EXPERIMENT DESIGN

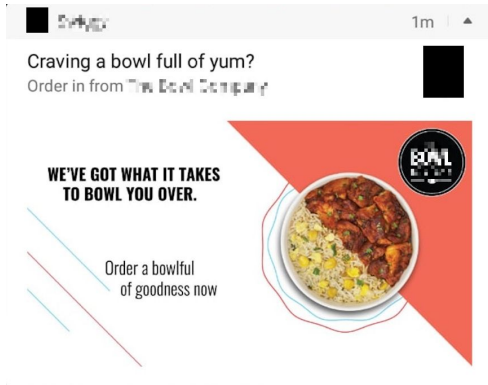
The data for this paper come from a food delivery mobile app platform in India. The platform serves multiple cities in India and is one of the largest players in the food delivery business. This firm partners with more than 50,000 restaurants and consumers can order meals from any of the partner restaurants that service their location through their mobile app. The firm regularly sends its customers in-app mobile push notifications informing them of offers, newly listed restaurants etc. The majority of offers advertised through push notifications are applicable on the entire platform or several restaurants on the platform. Some of these offers also inform customers of offers for specific restaurants on the platform. Consumers can click on these push notifications and directly go to the relevant page on the app to take advantage of the offer, or go to the app independently and select a restaurant to order from and then redeem the offer on their order.

As described in the introduction, to demonstrate the existence of the ‘discount spotlighting’ effect and to more broadly understand the effects of price vs brand advertising, I design an experiment to create exogenous variation in advertising intensity, price information content in advertising and prices for a focal restaurant on the platform. For the experiment, a focal restaurant on the platform was chosen to be advertised to consumers through in-app mobile push notifications. A sample of the firm’s customers were randomly assigned to receive these push notifications with a frequency of 0, 1, 2, 3 or 4 notifications a week for a period of four weeks. A discount was made available to customers as a percentage off of their total meal value. One of 5 discount levels: 0%, 10%, 20%, 30% or 40% was assigned randomly to consumers in the experiment sample. This discount was valid for all purchases¹ from the restaurant through the entire experiment period. The availability of the assigned discount was also independent of whether the consumer received any ads or not. Anyone

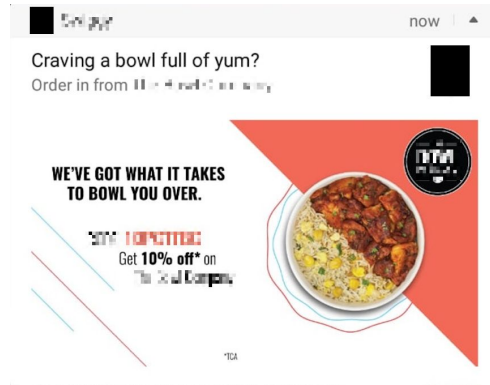
1. Upto a cap of four purchases a week and a maximum discount amount of Rs. 100 per order. Robustness checks with controls for ‘hitting the cap’ in terms of maximum number of orders per week or the maximum discount amount were done to ensure that these events don’t affect the main findings

who ordered from the focal restaurant also received free delivery on their order as a baseline offering, regardless of the discount level assigned. Customers in the experiment sample were also randomly assigned to one of two ad type conditions: Brand Ad and Discount Ad². The brand ad does not contain any information about the available discount. The price ad highlights the available discount percentage in block letters. Apart from information on the available discount, the remaining information content about the focal brand across the different ad types is the same. This helps us measure the causal effect of the mere inclusion of price information in advertising. The actual creatives for the different types of ads are shown in Figure 2.1. Figure 2.2 shows an actual screenshot of a phone to demonstrate how the notifications show up on a phone screen. Note that identifying information of the app and the focal restaurant have been removed from the ad creatives and different screenshots.


2. An additional creative (intermediate ad) that only mentioned that there was a discount available without mentioning the exact percentage discount and without highlighting it was also used. The effects of this ad were very similar to the brand ad and hence I do not show this separately in the results




(a) Brand Ad




(b) Discount Ad 10%



(c) Discount Ad 20%




(d) Discount Ad 30%

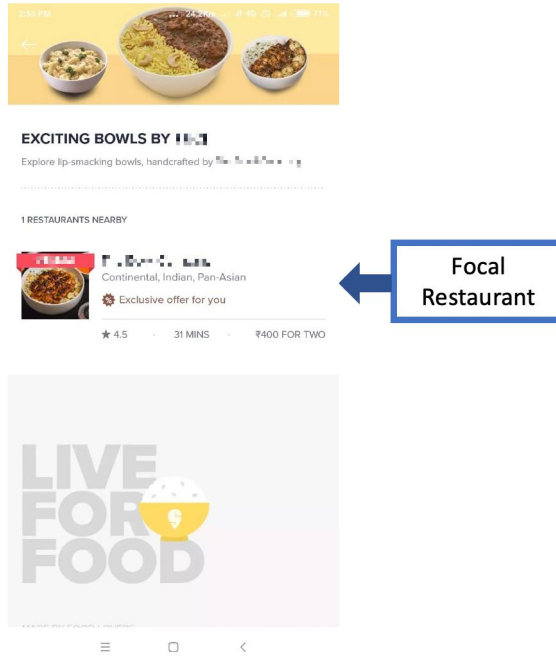


(e) Discount Ad 40%

Figure 2.1: Different Ad Creatives



(a) Push notification pop up on phone screen



(b) Landing Page after clicking the push notification

Figure 2.2: Screenshot of ad push notification on an Android phone and the landing page after clicking on the ad

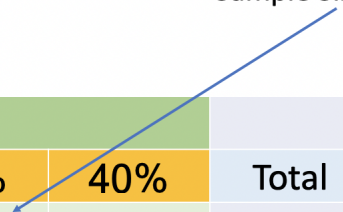
Randomization of discount level, ad frequency and ad type was done at the customer (app login ID) level and these were randomized independently of each other. One exception was that individuals who were assigned 0% discount were only assigned the Brand Ad condition (since they couldn't be given false information of a discount that was not available to them). To clarify the availability of the discount once again, the discount assigned to the consumer was available to them regardless of whether they received an ad or which type of ad they received. Consumers did not have to put in a coupon code to redeem the offer - the offer was automatically applied to any 'cart' that was created by the consumer for the focal restaurant.

Figure 2.3 summarizes the sample sizes in each cell of the experiment. A total of 149,823 individuals were included in the experiment.

The individuals in the experiment sample were chosen such that:

- They have an android device (as push notification delivery can only be tracked on

Sample Sizes



	Discount Level					
Type of Ad	0%	10%	20%	30%	40%	Total
No Ads	5,100	7,182	7,385	7,171	7,377	34,215
Brand Ads	24,227	11,359	11,345	11,485	11,431	69,847
Discount Ads	NA	11,519	11,603	11,369	11,270	45,761
Total	29,327	30,060	30,333	30,025	30,078	149,823

Figure 2.3: Experiment Design

android and not on iOS)

- They have made atleast one order on the app in the three months preceding the start of the experiment
- They were reachable by push notification during the week before the experiment
- They have had atleast three app sessions (i.e. they have opened the app atleast three distinct times separated by a gap of atleast 90 minutes) in the month preceding the start of the experiment in which the focal restaurant has appeared in their restaurant listings. This is to ensure that they live or work in the areas serviceable by the focal restaurant (only serviceable restaurants show up in the listings).

To make a purchase from a restaurant, the consumer has to first go to the restaurant menu page. A consumer can reach the menu page of the restaurant in the following ways:

- Click on the push notification ad, following which the app opens on the phone and the consumer is taken to a landing page with a link to the restaurant menu page

- Open the app independent of the ad, and click on the focal restaurant on the restaurant listings page
- Search for a cuisine or the restaurant name, following which the focal restaurant shows up as a listing in the search results, which the customer can then click

At the menu page, the different dishes available in the restaurant are displayed along with their individual prices (without discount). The consumer can add dishes that they are interested in purchasing to their cart. At cart, the final price of the meal after discount is displayed to the consumer. She then makes the decision of whether to purchase the meal at the displayed price or not. If she decides to purchase, she can enter her exact delivery location, pay and finish ordering. The different stages of the purchase funnel are diagrammatically represented in Figure 2.4.

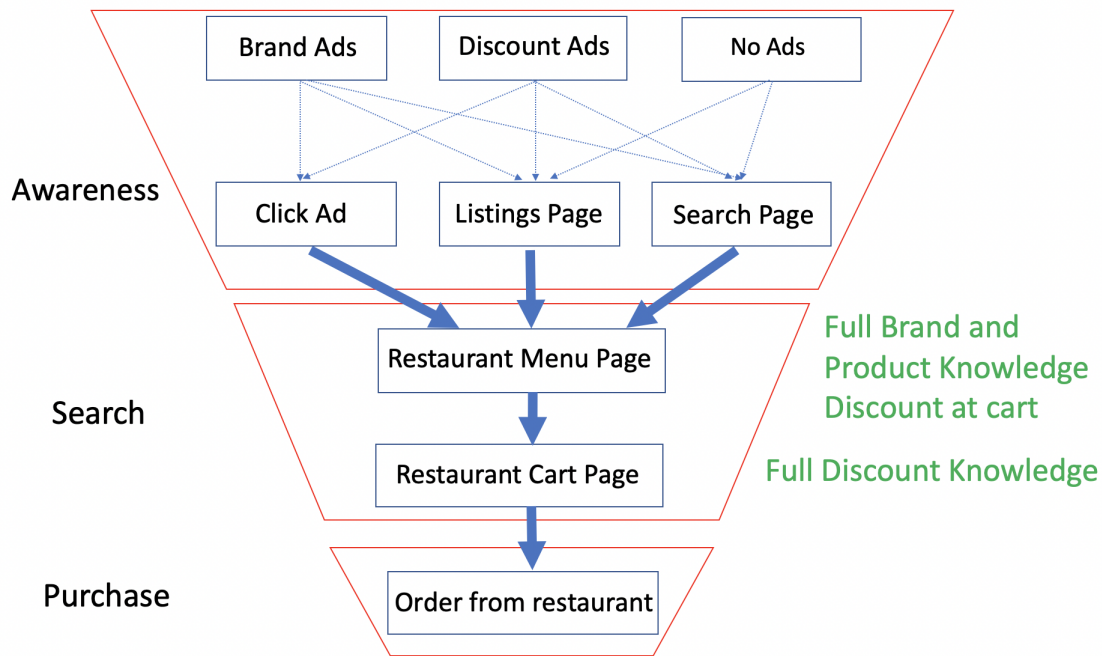
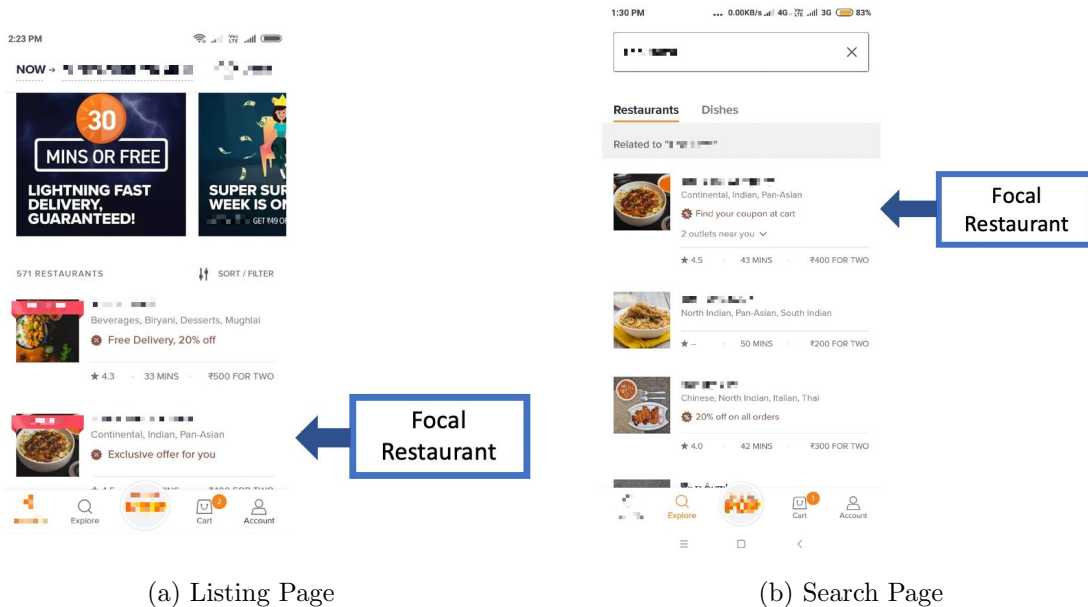


Figure 2.4: Steps in the purchase funnel

Screenshots of the different pages along the customer purchase funnel described above are shown in Figures 2.2, 2.5 and 2.6.

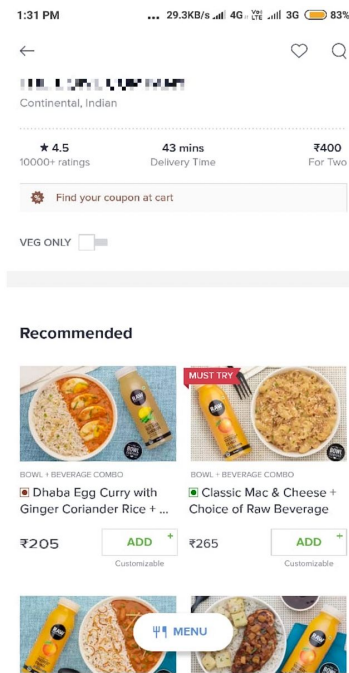
It is important to note that the exact discount percentage available to each customer



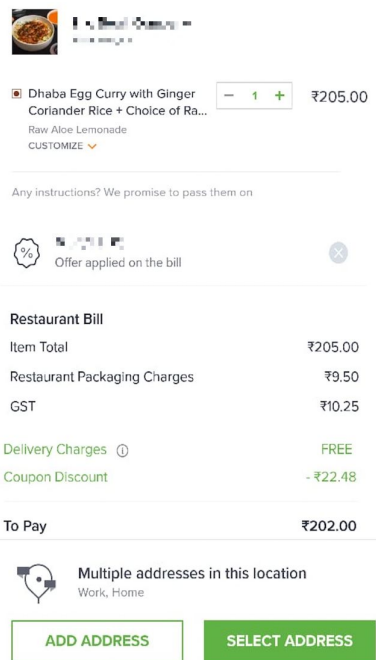
(a) Listing Page

(b) Search Page

Figure 2.5: Listing and Search pages



(a) Menu Page of Focal Restaurant



(b) Cart Page for Focal Restaurant with assigned discount auto-applied

Figure 2.6: Menu and Cart pages

is either disclosed to the consumer through the price ad, or is revealed after the consumer visits the cart page. Thus, until the consumer visits the cart page, the differences along the purchase funnel arise due to the different levels of discount information disclosed to her through the different types of ads. However, upon reaching the cart, all consumers see full information about the discount regardless of the type of ad they were assigned to. This setup allows us to disentangle the ‘discount spotlighting’ effect from the ‘brand advertising effect’ and the ‘pure promotion effect’. At this stage, since brand knowledge and knowledge of the discount are the same across all consumers, and the only difference between them is the type of ad they saw, any differences in conversion to purchase are a result of the ‘discount spotlighting effect’.³

3. There could also be differences as a result of increased price sensitivity due to price advertising or a reference price effect from the consumer seeing a discount that is higher or lower than their expectations (for brand ads). I rule out these alternate explanations later

2.1 Data

I observe a rich set of historical data as well as data generated during the experiment for each consumer - their purchases, their visits to different pages along the purchase funnel, including visits to menu and cart pages of any restaurant. I observe the total meal value that the customer adds to their cart and the final price after applying the available discount. I also observe whether each push notification was sent, received, and clicked on (both for historical and experimental push notifications).

When a push notification is sent, it may not be received by the consumer due to various reasons - she has turned off notifications or uninstalled the app, she is out of reach of the network or in some cases the phone suppresses push notifications in case battery level is below 15%. In case a notification is not received by the consumer, I can observe whether that event is due to either a ‘send error’ i.e. the app has been uninstalled or push notifications have been turned off, or a ‘receive error’ i.e. the phone is out of reach of network or push notifications have been temporarily suppressed by the phone. About 44.3% of the sample received atleast one fewer notification than what they were assigned. Most of these instances are due to a single notification not being delivered. Table 2.1 shows the share of customers that received n fewer notifications than assigned for all values of n upto 16. 13.3% out of the total 44.3% received fewer notifications than they were assigned due to a ‘send error’ i.e. they turned off notifications or uninstalled the app atleast for some time during the experiment. The remaining 31% received fewer notifications due to network or battery issues.

I deal with potential self selection concerns due to people receiving a different number of notifications than what they were randomly assigned to receive in two ways: (1) For some parts of the analysis, I just use the assignment of customers to different types of ads without using the frequency. Only 26 people in the entire sample received zero notifications when they were not assigned to zero ad frequency. Thus, virtually the entire sample received atleast one ad, if they were assigned to receive any ads. Thus, the effect of having being assigned to a particular ad type and receiving atleast one ad is correctly captured. (2) For

Table 2.1: Difference between total number of notifications randomly assigned and actually received

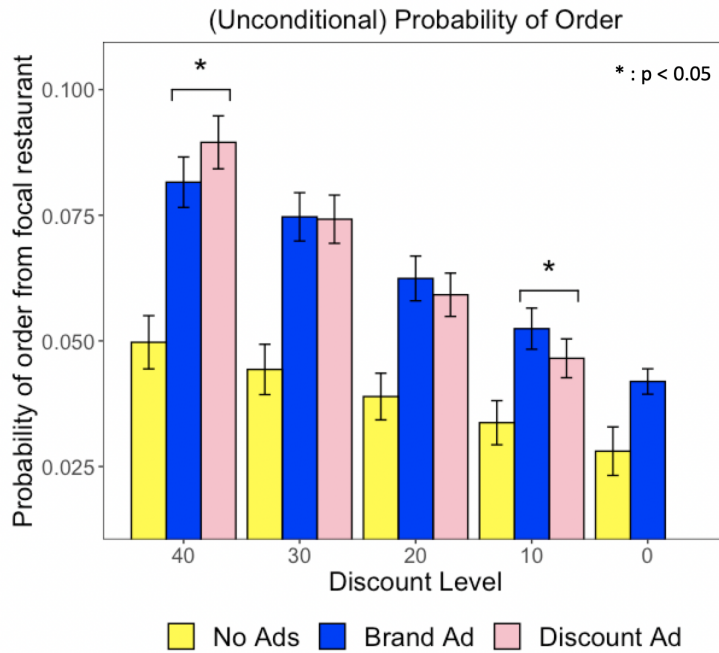
Difference	% of customers
1	21.46%
2	8.81%
3	4.35%
4	2.31%
5	1.53%
6	1.53%
7	1.23%
8	0.58%
9	0.62%
10	0.55%
11	0.67%
12	0.15%
13	0.15%
14	0.16%
15	0.21%
16	0.001%
Atleast 1	44.3%

parts of the analysis that uses a goodwill stock model of advertising, which uses ad frequency, I do a robustness check by restricting the sample to just the people who received exactly the number of ads they were assigned experimentally. Another robustness check is done by replacing the goodwill stock of advertising by a dummy variable indicating whether an ad has been seen on or before by the consumer to measure the effect of having seen at least one ad of the assigned type. All results qualitatively carry through. The results from the experiment are described in the following chapters.

CHAPTER 3

EFFECT ON DEMAND

In this chapter, I study the effects of the different types of ads on demand for the focal restaurant. Figure 3.1 shows the probability of a customer placing at least one order with the focal restaurant under the different ad conditions.



Error bars report 95% confidence intervals

Figure 3.1: Effect of the different types of ads on demand for the focal restaurant

We see that all ads raise demand at all discount levels. Ad effects are bigger at higher discount levels, i.e. advertising shifts the demand curve out and also makes it more elastic (consistent with Erdem et al., 2008). Comparing the different types of ads, we see that at low discounts (10%), brand advertising leads to highest demand and at high discounts (40%), price advertising results in highest demand. The demand curve is thus more elastic under price advertising relative to brand advertising (i.e. the difference in demand between 10% and 40% discount is higher under price advertising relative to brand advertising). Among

the tested discount levels in the experiment, the optimal discounts¹ are: 20% under no advertising, 30% under brand advertising and 40% under price advertising, with 40% under price advertising being the best overall. This illustrates that when setting prices, managers should first measure elasticities under the different advertising options and then jointly optimize over both price and type of advertising. Also, the finding that brand advertising leads to higher overall demand at low discounts highlights that caution should be exercised when applying the conventional wisdom in digital advertising that suggests that any available discount should be highlighted in digital ads. We will explore the mechanism behind this finding in the next chapter.

Figure 3.1 does not utilize the variation in ad frequency. Next, I analyze the full panel with data at the individual-day level and utilize the variation in ad frequency to measure ad effects on demand using the following model:

$$y_{it} = \alpha + \gamma_1 ad_{it} + \gamma_2 ad_{it}^{price} + \beta discount_i + \nu_1(ad_{it}) * (discount_i) + \nu_2(ad_{it}^{price}) * (discount_i) + \epsilon_{it} \quad (3.1)$$

where

- y_{it} is an indicator variable that is 1 if individual i ordered from the focal restaurant on day t
- ad_{it} is the goodwill stock of advertising for individual i on day t . Goodwill stock is defined as $ad_{it} = \sum_{\tau=0}^t \delta^{t-\tau} a_{it}$ where a_{it} is an indicator variable that is 1 if individual i received an experimental notification on day t of the experiment. δ is the advertising carryover factor. Following Shapiro et al., 2018, I use a grid search to fix the value of δ . Model (3.1) is estimated using OLS repeatedly with different values of δ starting

1. Since costs are not known, these are revenue maximizing prices (after accounting for the discount amount)

from 0 to 1 with increments of 0.01. The value that returns the best value of R^2 is used. Following this process, δ is fixed at 0.24.

- ad_{it}^{price} is the goodwill stock of price advertising for individual i on day t . The coefficient on this variables will indicate the difference in the ad effect between price advertising and brand advertising. To get the total price ad effect, the coefficients on ad_{it} and ad_{it}^{price} must be added. The interpretation of coefficients is similar for the interaction terms of discount and advertising ²
- $discount_i$ is the percentage discount that is assigned to individual i i.e. 0,10,20,30 or 40

Estimation results using a probit specification as well as OLS for model (3.1) are shown in Table 3.1 under specification I and II respectively. For probit, the specification is changed to

$$y_{it}^* = \alpha + \gamma_1 ad_{it} + \gamma_2 ad_{it}^{price} + \beta discount_i + \nu_1(ad_{it}) * (discount_i) + \nu_2(ad_{it}^{price}) * (discount_i) + \epsilon_{it} \quad (3.2)$$

and

$$y_{it} = 1 \text{ if } y_{it}^* > 0 \text{ and } 0 \text{ otherwise; } \epsilon_{it} \sim N(0, 1) \quad (3.3)$$

We see from both specifications that ads have a positive effect on demand. Further, the intercept term for price advertising is negative indicating that at low discounts, price advertising performs worse than brand advertising. However, the slope of demand (interaction between discount and advertising) is higher for price advertising relative to brand

2. Similar terms are also included for intermediate ads. Brand ads and intermediate ads are found to have similar effects. This holds true for all results in the remaining chapters of this dissertation as well. We are primarily interested in interpreting the differences between price and brand ads. Thus, in the interest of brevity, results for intermediate ads are not reported separately and the terms in equations are not shown, although they are used in estimation

Table 3.1: Effect of Ads on Demand for the Focal Restaurant

	DV: Ordered from restaurant				DV: Meal value ordered	
	Probit (I)		OLS (II)		OLS(III)	
	Estimate	SE	Estimate	SE	Estimate	SE
Intercept	-2.896***	0.009				
<i>ad</i>	0.131***	0.014	0.001***	0.0001	0.314***	0.054
<i>ad^{price}</i>	-0.069**	0.028	-0.001***	0.0003	-0.249*	0.113
<i>discount</i>	0.007***	0.0003	0.00006***	0.000003	0.018***	0.001
<i>discount * ad</i>	-0.0003	0.0005	0.000019**	0.000007	0.0058*	0.0026
<i>discount * ad^{price}</i>	0.0021*	0.001	0.000034*	0.00001	0.0082(.)	0.0047
Day FE			Yes		Yes	
No. Obs	5,474,448		5,474,448		5,474,448	

Significance Codes: *** $p \leq 0.001$, ** $p \leq 0.01$, * $p \leq 0.05$, (.) $p \leq 0.1$

Std. Errors clustered at the individual level

advertising, indicating that as discount level increases, price advertising does better than brand advertising. So demand at high discount levels is higher under price advertising. The OLS specification also tells us that the slope of demand under brand advertising is higher than that under no advertising. Thus, we confirm the insights from Figure 3.1. Advertising has a positive effect on demand and it increases the elasticity of demand. Price advertising performs worse than brand advertising at low discount levels but better at high discount levels i.e. elasticity of demand is higher under price advertising.

Another specification where the dependent variable is changed to the total meal value (order value before discount) is also shown under specification (III), leading to the same conclusions. Under this specification, we capture the effect of discounts and advertising on increasing the probability of ordering from the focal restaurant as well as the effect on increases in meal value conditional on ordering. It is interesting to examine whether the way in which advertising and discounts have an effect on total demand is through increasing the number of orders or also through increasing the meal value that people add to their carts when ordering. In other words, do people add more items to their cart and increase the size of the order under high discount and advertising conditions. To examine this, I run a regression with meal value added to cart(i.e. monetary value of the order before applying

the discount) as the DV on discount level and ad stock, for the individual-day combinations when a cart is created for the focal restaurant. I also control for the pre-experiment average order value of each individual to control for self selection of individuals that differ in pre-experiment order size under the different advertising and discount conditions. The results of this regression are reported in Table 3.2.

Table 3.2: Meal value conditional on preparing a cart for the focal restaurant

	Estimate	SE
Discount	-0.29	0.56
Ad Stock	-2.36	2.95
Pre-expt. avg. order value	0.42***	0.03

Std. Errors clustered at the individual level

We see that after accounting for pre-experiment average order value, there seems to be no effect of discounts or advertising on increasing meal size size conditional on placing at least one item in the cart for the focal restaurant. Thus, the way that advertising and discounts seem to work in this context is through the extensive margin i.e. getting more people to place an order or place orders more frequently. There is no effect on meal value conditional on an order being placed. Thus, for the remainder of the paper, I will only focus on the probability of placing an order and not on the size of the order.

Now that we have examined the effects on demand, we will next look into the different stages of the purchase funnel to investigate the reasons for the differences in demand arising from price advertising and brand advertising. Do the differences in demand arise from consumers searching differently, or do they arise from differences in consumers converting to purchase at different rates once full information about price has been revealed to them?

CHAPTER 4

EFFECT ON SEARCH

One of the primary mechanisms by which advertising is thought to affect demand is through its informative effect on search (Stigler, 1961; Ozga, 1960) . In this chapter, we will focus on the arrival rates of consumers at the upper stages of the purchase funnel. The main finding here is that the rate at which consumers visit the menu page of the focal restaurant differs with discount level for price advertising, but not with brand advertising. Consumers self-select themselves into visiting the restaurant menu page in response to different types of ads according to their knowledge or expectation of available discount. This demonstrates the primary purpose of conveying discount information to consumers using price advertising - managers do this hoping that many consumers check out the product (i.e. search for it), being attracted by the presence of a high discount. If discount information is not conveyed through ads, consumers have to make the search decision based on their expectations of price. In what follows, I will discuss this in more detail.

If price information in advertising affects how consumers search, we should expect to see differences in rates at which consumers who were exposed to different ads with varying price information arrive at the menu page of the restaurant.

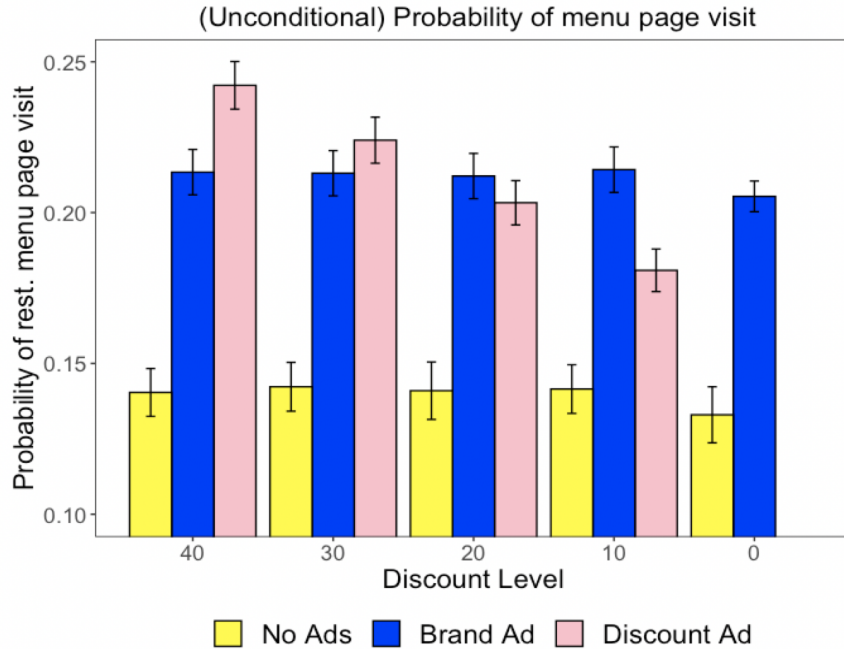
Consumers who received brand ads, that do not mention a discount, presumably make the decision to search based on their expectation of available discounts. Most notifications that the platform has sent consumers historically convey information on discounts. Thus, consumers may expect a non-zero discount conditional on receiving a notification from the platform, even if the notification does not explicitly mention one. The consumer will pay a search cost and visit the menu page of the restaurant only if the expected discount justifies this decision. On the other hand, if consumers are explicitly told the exact available discount, they can make the search decision under better information i.e. they will pay the search cost of going to the menu page of the restaurant if they think that the available discount justifies the cost of search.

According to the above hypothesis, if the expectation of discount conditional on receiving a notification about a restaurant is high (above 10%), then brand advertising may lead to higher search relative to price advertising at lower discount levels. This is consistent with the notion of uninformative advertising as an invitation to search (Mayzlin and Shin, 2011). Conversely, if the expectation of discount is low (below 10%), then price advertising should lead to more search at all discount levels in the experiment.

Based on historical push notifications sent to each customer in the experiment sample, I create a measure of expected discount conditional on receiving a notification for each customer (which is simply the average of all discount amounts mentioned in previous notifications received). The average expected discount conditional on receiving a notification is found to be 23%. Based on this number and the above hypothesis of search based on discount expectations, we should predict that brand advertising leads to higher search rates at 0,10% and 20% discount levels and price advertising leads to higher search rates at 30% and 40%.

Figure 4.1 shows the probability of an individual visiting the menu page of the focal restaurant during the experiment under the different ad and discount conditions.

First, we see that ads have a positive effect on search probability at all discount levels. We also see that the probability of search i.e. visiting the restaurant menu page does not vary with discount level under the ‘No Ad’ and ‘Brand Ad’ conditions. This makes intuitive sense since people cannot respond to information that they haven’t been given. However, if given price information through price ads, we see that consumers respond to the different discount levels by searching more if there is a high discount available and less if a low discount is available. However, this is not necessarily a good thing for the brand. At low discount levels (10%), we see that a lower number of people search for the brand. If the brand wants to maximize the number of people searching for it under the availability of a 10% discount and at current consumer beliefs (23% expected discount conditional on receiving an ad), it is better off using brand ads. The higher level of search at low discounts with brand advertising



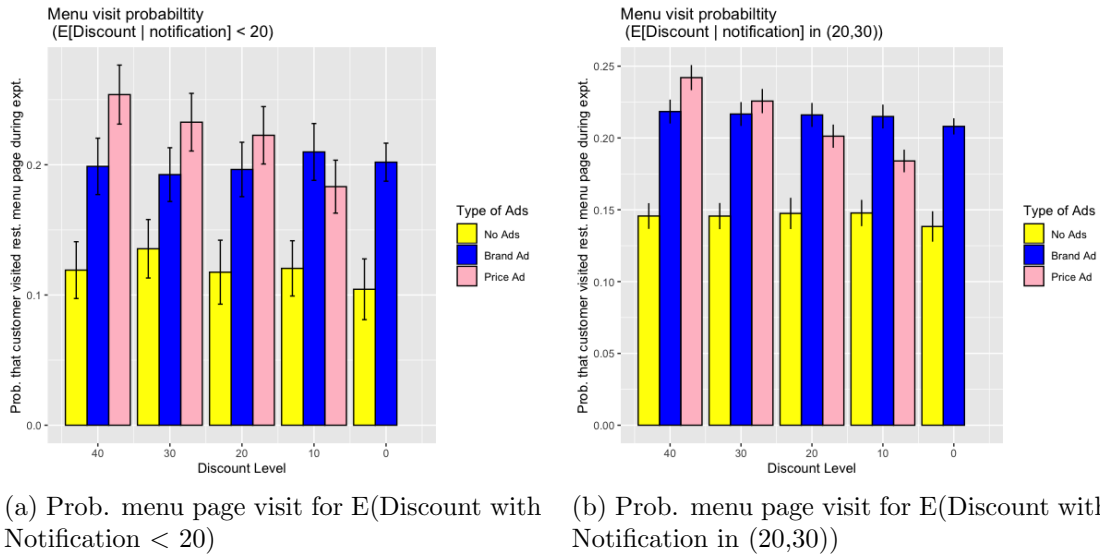
Error bars report 95% confidence intervals

Figure 4.1: Effect of the different types of ads on probability of visiting the focal restaurant menu page

also translates to higher overall demand as we saw in the previous chapter. This is empirical evidence consistent with the notion of uninformative advertising as an invitation to search (Mayzlin and Shin, 2011). Conventional wisdom of mentioning any available discount in ads is proved to be wrong in this case (and is generalizable to similar high discount contexts eg. Groupon[®]). However, at high discount levels, the informative effect of price advertising works in favor of the firm i.e. leads to higher search and demand.

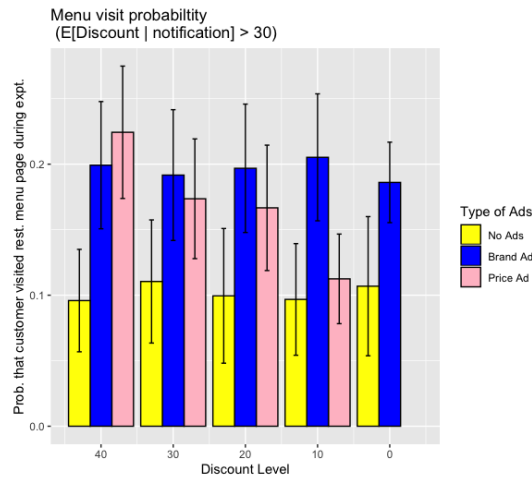
We see that the search probabilities for brand ads are between those for the 20% and 30% price ads. This is consistent with the notion of search based on expected discounts and the fact that the average expected discount conditional on receiving a notification is found to be 23%. In order to delve deeper into this, I split the sample of individuals according to the measure for expected discounts and look at the search patterns for different sub-samples. Figure 4.2 shows the probability of visiting the restaurant menu page for individuals with expectations of discount split into three bins - less than 20, between 20 and 30; and greater

than 30. This figure demonstrates that customers who receive brand ads make the decision to search based on their expectations of discount



(a) Prob. menu page visit for $E(\text{Discount with Notification} < 20)$

(b) Prob. menu page visit for $E(\text{Discount with Notification in } (20,30))$



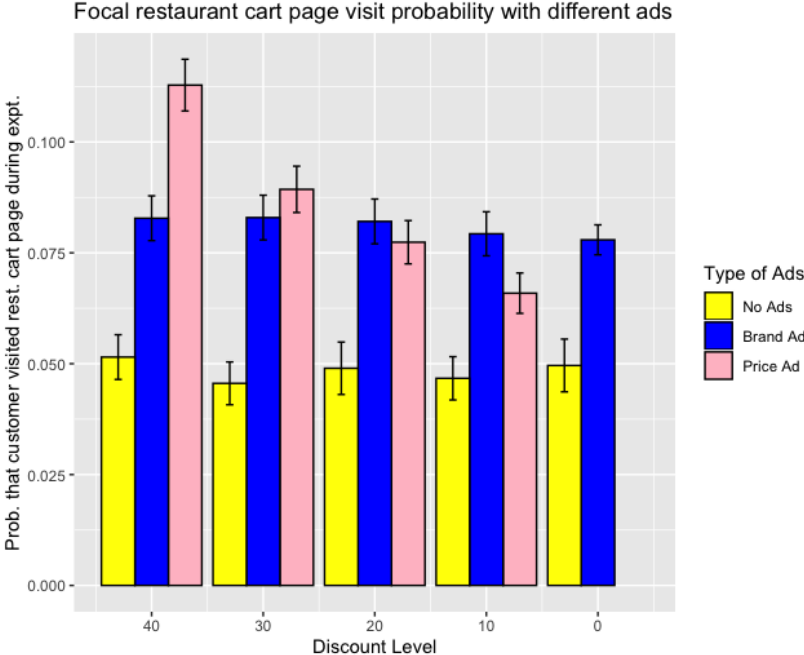
(c) Prob. menu page visit for $E(\text{Discount with Notification} > 30)$

Figure 4.2: Effect of different types of ads on the probability of visiting the focal restaurant menu page on sub-samples of consumers split by expectations of discount conditional on receiving a notification

We see that the search probabilities for brand ads are between those for the 10% and 20% price ads for consumers whose discount expectations are less than 20. Similarly, the search probabilities for brand ads are between those for the 30% and 40% price ads for consumers whose discount expectations are greater than 30. Most of the sample falls in the ‘between

20 and 30' range and their search probabilities for brand ads are between those 20% and 30% price ads. This demonstrates that people indeed search according to their expectations of discount if they are not informed of specific discount level through price advertising. The notion of a 'brand ad' in this context might be different from that in the literature, where the assumption is that a brand ad does not evoke any expectation of a non-zero discount.

Next, we look at the second stage of the purchase funnel, which is the cart page. Figure 4.3 shows the probability that an individual visits the focal restaurant cart page after adding items to the cart from the menu page.



Error bars report 95% confidence intervals

Figure 4.3: Effect of the different types of ads on probability of visiting the focal restaurant cart page

We see a similar pattern as we saw for the menu page. Since price information is only revealed after visiting the cart page, the probability of visiting the cart page does not change vary with discount level under the 'No Ad' and 'Brand Ad' conditions. However, it does vary with discount level under the 'Price Ad' condition in the way that we would expect.

4.1 Self-selection of consumers into search

Since consumers are able to make the search (menu page visit) decision with more information under price advertising, this means that they are able to self select into search according to their knowledge of the available discount and their sensitivity to discounts. We might expect that individuals who are highly discount seeking or most price sensitive make the decision to visit the menu page only under the 40% condition. On the other hand, the people who make the decision to visit the menu page when told there is only a 10% discount should be relatively less discount seeking. Since I observe past purchase behavior of all individuals, I can characterize the people who visited the menu page under the different ad and discount conditions - and test the hypothesis that the group of individuals who visited the menu page when told that there is a high discount (or expect a high discount when not told explicitly) are on average more discount seeking than the group on individuals who searched when told that there is a low discount.

Based on pre-experiment purchases I create the following proxies for consumer ‘discount-seeking’ or price sensitivity:

- Fraction of orders purchased on discount
- Average discount amount used per order conditional on having used a discount
- Average discount amount used per order (unconditional)
- Average ‘cost for two’ descriptor for restaurants purchased from. Each restaurant on the platform contains a descriptor called ‘Cost for two’ which indicates the price of an average meal for two people, at that restaurant. This is a characteristic that is provided by the restaurant owner at the time of signing-up with the firm. This information appears on the restaurant listings and menu pages. We can interpret this quantity as indicating whether a consumer orders from relatively more expensive or cheap restaurants.

Apart from these proxies for discount-seeking or price sensitivity, we might also expect that individuals who are not familiar with the focal restaurant, i.e. those who have not made a single purchase from the focal restaurant before the experiment, are more likely to respond only to high discounts. Thus I create a dummy variable that indicates whether the consumer is a ‘previous restaurant customer’ i.e. has ordered atleast once from the focal restaurant before the start of the experiment.

Figures 4.4 and 4.5 show the mean of these customer ‘characteristics’ for individuals who chose to visit the focal restaurant menu page (i.e. self selected into search) during the experiment.

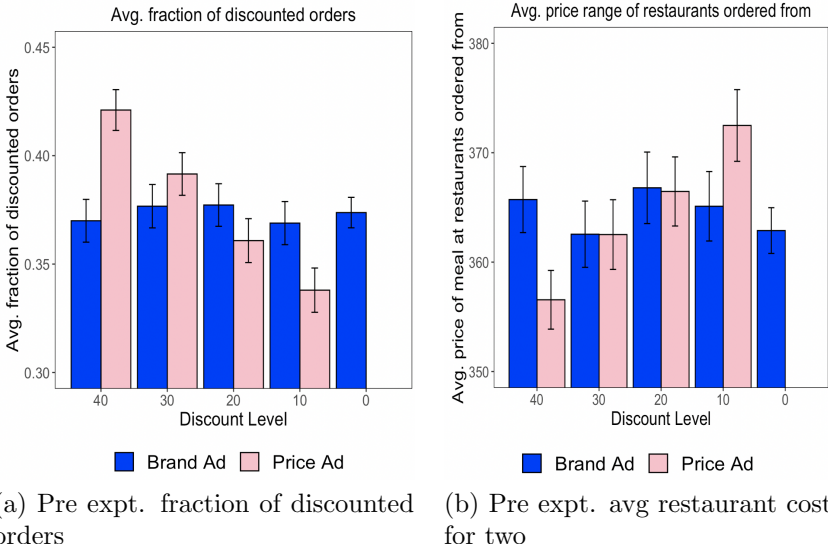


Figure 4.4: Pre experiment customer ‘characteristics’ for individuals who chose to visit the focal restaurant menu page under different discount and ad conditions (1)

From Figure 4.4 we see that individuals who responded to the 40% discount price ad are those who have made a large share of their previous purchases using a discount i.e. they are more discount seeking. Individuals who responded to the 10% discount price ad are those who have made a relatively smaller share of their previous purchases using discounts. We see a similar pattern when we examine the ‘cost for two’ of previous restaurants ordered from. Individuals who searched in the 10% price ad condition are people who order from relatively more expensive restaurants and the individuals who searched in the 40% price ad condition

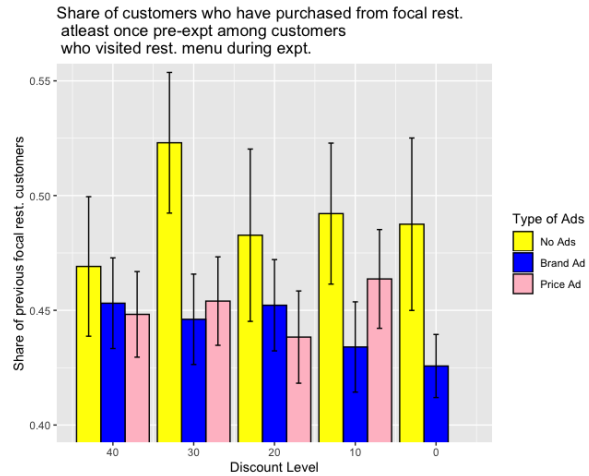
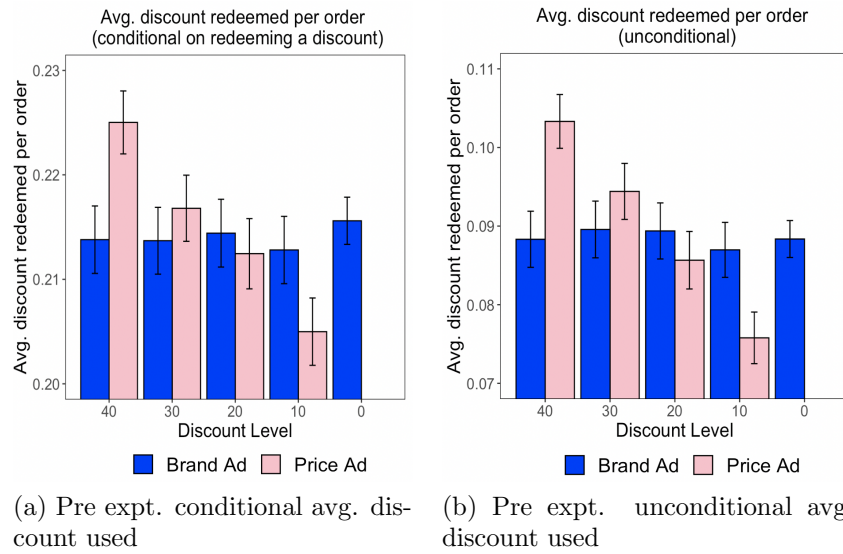


Figure 4.5: Pre experiment customer ‘characteristics’ for individuals who chose to visit the focal restaurant menu under different discount and ad conditions(2)

are those who order from relatively cheaper restaurants. Individuals who visited the menu page in the ‘No Ad’ and ‘Brand Ad’ conditions are pretty similar to each other in terms of these characteristics and they also look similar to the individuals who responded to the 20% and 30% price ads.

Individuals who received no ads also see a message saying ‘Exclusive Offer for you’ on the restaurant listings page, which might have created an expectation of discount between

20% and 30% for them as well. Restaurants which offer discounts on the platform usually offer discounts in this range, so such an expectation is rational. This would explain why the ‘No ad’ responders look similar to those who responded to the the brand ads in terms of their pre experiment characteristics. The higher level of search under brand advertising is then purely due to the informative effect of brand advertising relative to no advertising. In terms of characteristics indicating price sensitivity, the people who responded to brand ads seem similar to those who responded without ads.

From Figure 4.5 (a) and (b), we see that that individuals who responded to the 40% discount price ad are those who have used high discount amounts on their pre-experiment orders i.e. they are more discount seeking. Individuals who responded to the 10% discount price ad are those who have used lower discount amounts for their pre-experiment purchases. From, Figure 4.5 (c), we can see that all the different forms of advertising lead to more ‘new to restaurant’ customers visiting the menu page of the focal restaurant. This is in-line with the informative effect of advertising i.e. it informs new consumers of the existence of this restaurant. However, the different ad types or the different discount levels do not attract new customers at different rates.

CHAPTER 5

EFFECT ON CONVERSION TO PURCHASE CONDITIONAL ON SEARCH

This is the chapter that drives the main result of this dissertation. By comparing rates of conversion to purchase among consumers who received different ads (that were randomly assigned to them) and are at the cart stage where they have full product information including the available discount, we can test for the residual impact of discount information in advertising, after differences in consumer knowledge of the available discount have been eliminated. Self-selection into visiting the cart page makes this comparison problematic (as we are no longer able to directly compare identical sets of consumers) but I address this using both traditional econometrics methods and modern machine learning methods. After accounting for self-selection using these different methods (i.e. comparing arguably identical sets of consumers who were randomly assigned different ads), I show that consumers who receive price advertising convert to purchase at lower rates compared to those who received brand advertising or no advertising. Since we are comparing arguably identical consumers who have the same knowledge of the brand and the available discount, and they only differ in whether they were randomly assigned to receive price ads or brand ads, the difference in conversion rates to purchase must be due to the ‘discount spotlighting’ effect. I also rule out alternate explanations like increased price sensitivity due to price advertising or reference price effects. Having said this, we must also remember that the overall revenue maximizing strategy for the restaurant was a 40% discount with price advertising. The negative effect on brand preference was offset by the fact that many people visited the restaurant menu page attracted by the high discount and enough of them purchased to make this an optimal strategy. This explains the prevalent use of price advertising combined with high discounts, despite possible negative effects brought up in this paper and previous literature.

In what follows, I go into more detail.

Since the different ad formats lead to different types of consumers arriving at the cart stage through self selection into search, conditional conversion rates at cart from the different ads may not be directly comparable. This is the classic self-selection issue that many econometricians have faced before (Heckman, 1979; Greene, 2000).

I approach this issue in two ways:

1. Use a matching on observables approach through a non-parametric Causal Forest setup, using a high dimensional set of consumer characteristics constructed from pre-experiment historical data
2. Explicitly model the selection and conversion steps as separate stages of the consumer decision process and allow for correlation in the error terms across these steps

The first approach relies on the *unconfoundedness* assumption (Rubin, 1990) to recover the difference in treatment effects of brand ads and price ads. The assumption is that after controlling for observable characteristics, biases in comparisons between individuals who were treated with brand and price ads and arrived at cart are removed (i.e. we are comparing arguably identical individuals who have been randomly assigned the different ad types), thus allowing for a causal interpretation of those adjusted differences. However, there may be unobservables that are not captured through the set of characteristics created from historical data. For example, the proxies for price sensitivity may not capture some aspect of ‘true’ price sensitivity, which is a possible confound that is unobservable. Thus, I also include a second approach that allows for selection on both observables and unobservables, but relies on a distributional assumption about how the error terms between the different stages of the decision process are correlated.

5.1 Matching on observables using a causal forest

First, I use a matching on observables approach. One way to approach this would be propensity score matching (Imbens and Rubin, 2015). However this relies on imposing a parametric

specification in estimating the propensity score. Recent advances in machine learning (Wager and Athey, 2018) have made it possible to match on a high dimensional set of consumer characteristics in a non-parametric way while also allowing for arbitrary interaction effects between variables, using a causal forest. A causal forest avoids the necessity to impose a parametric specification and is computationally efficient, robust to model mis-specifications, and achieves desired consistency and asymptotic normality, and is thus a preferable approach.

I use a causal forest to estimate the difference in treatment effects between price and brand ads. The data used is the subset of data when individual i has created a cart for the focal restaurant. Further since we are interested specifically in measuring the difference in treatment effects between brand and price ads, I restrict attention to the subsample which have been assigned either brand or price ads and were assigned a positive discount. The individuals assigned to brand ads are taken as the control group and the individuals assigned to price ads are taken as the ‘treatment’ group. A dummy variable indicating conversion to order is the outcome variable. The list of ‘X’ variables that are used for matching include characteristics describing (1) previous purchase behavior on the platform: average order value, average discount redeemed, fraction of orders on discount, average price range of restaurants ordered from, total number of orders on platform, length of time active on platform, average number of orders a week, number of different restaurants ordered, number of different cuisines ordered (2) previous experience with the focal restaurant and similar restaurants: number of orders from previous restaurant, number of times visited menu page of focal restaurant, time since last order from restaurant, percentage of orders made from restaurants with same cuisine, total number of different cuisines ordered from (3) previous responses to notifications from the platform: fraction of previous notifications clicked, number of notifications received, average discount level in previous notifications, standard deviation of discount level in previous notifications and (4) previous search behavior on platform: average number of restaurants searched before order, number of restaurants searched with same cuisine as focal restaurant.

The quantity being estimated by causal forest (a Conditional Average Treatment Effect) is

$$(Pr[Purchase|Search, X_i, Discount_i, DiscountAd] - Pr[Purchase|Search, X_i, Discount_i, BrandAd])$$

and the output of the causal forest is the entire distribution of these differences for individuals in the experiment

The estimated forest reports an average treatment effect of -0.047 with standard error of 0.007. This means that the probability of conversion to purchase at cart for individuals assigned to price ads was on average 0.047 lower than individuals assigned to brand ads. Since the causal forest allows us to estimate heterogenous treatment effects and predict the estimated treatment effect for each individual in the sample, we explore this heterogeneity.

Figure 5.1 shows the the entire distribution of the difference in conversion rates between brand ads and price ads for the full sample of individuals who created a cart during the experiment.

We see that some individuals convert at higher rates with price ads, but for most individuals the conversion rate with brand ads is higher, 0.047 being the average difference in probability of conversion. Figure 5.2 shows the distribution of these effects at different discount levels.

We see at all discount levels, price ads perform worse than brand ads in terms of conversion rate. Also, differences between price ad and brand ad conversion rates increase with discount level i.e. the negative effect on brand preference is greater at high discount levels. Thus, a brand choosing to highlight a deep discount in its advertising could suffer a loss of brand equity.

Figure 5.3 shows the effects for new and previous customers of the focal restaurant.

We see that the difference in conversion rates is higher for new customers. Also, the

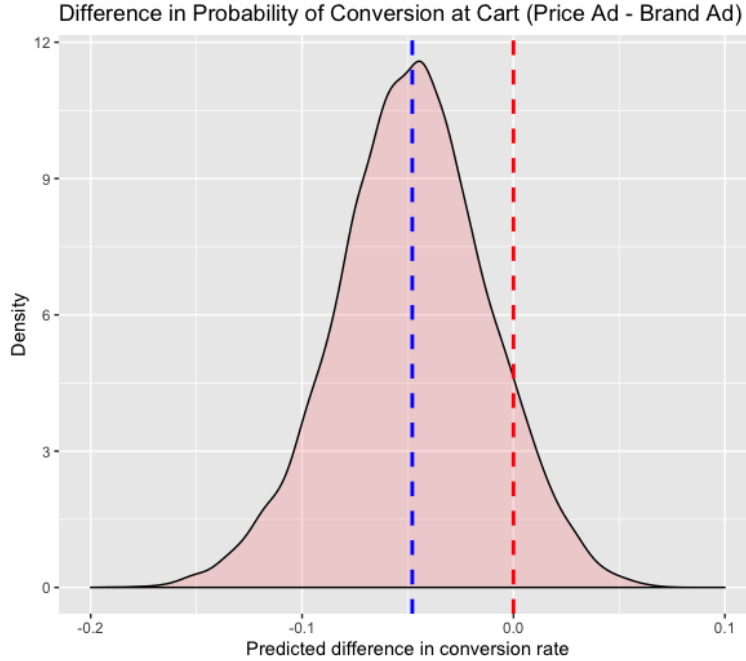


Figure 5.1: Distribution of differences in probability of conversion at cart after having been treated with a price ad and brand ad

difference completely vanishes for highly experienced consumers (more than 5 previous orders from focal restaurant). This is inline with our intuition as new customers do not have well-formed brand preference for the focal restaurant and hence are more likely to be influenced by the fact that it chose to highlight a discount in its advertising.

5.2 Modeling the consumer decision process to account for selection on unobservables and observables

The final approach to deal with self selection into visiting the cart page is to explicitly model the selection process. I specify a model where a consumer first makes the decision to visit the focal restaurant menu page based on an expectation of the final price to be paid depending on the type of ad received, her base brand preference for the restaurant, search cost and any residual advertising effect. The decision to visit the cart page from menu page is made based on the same factors but the effects of these factors can be different at this step. Finally, at

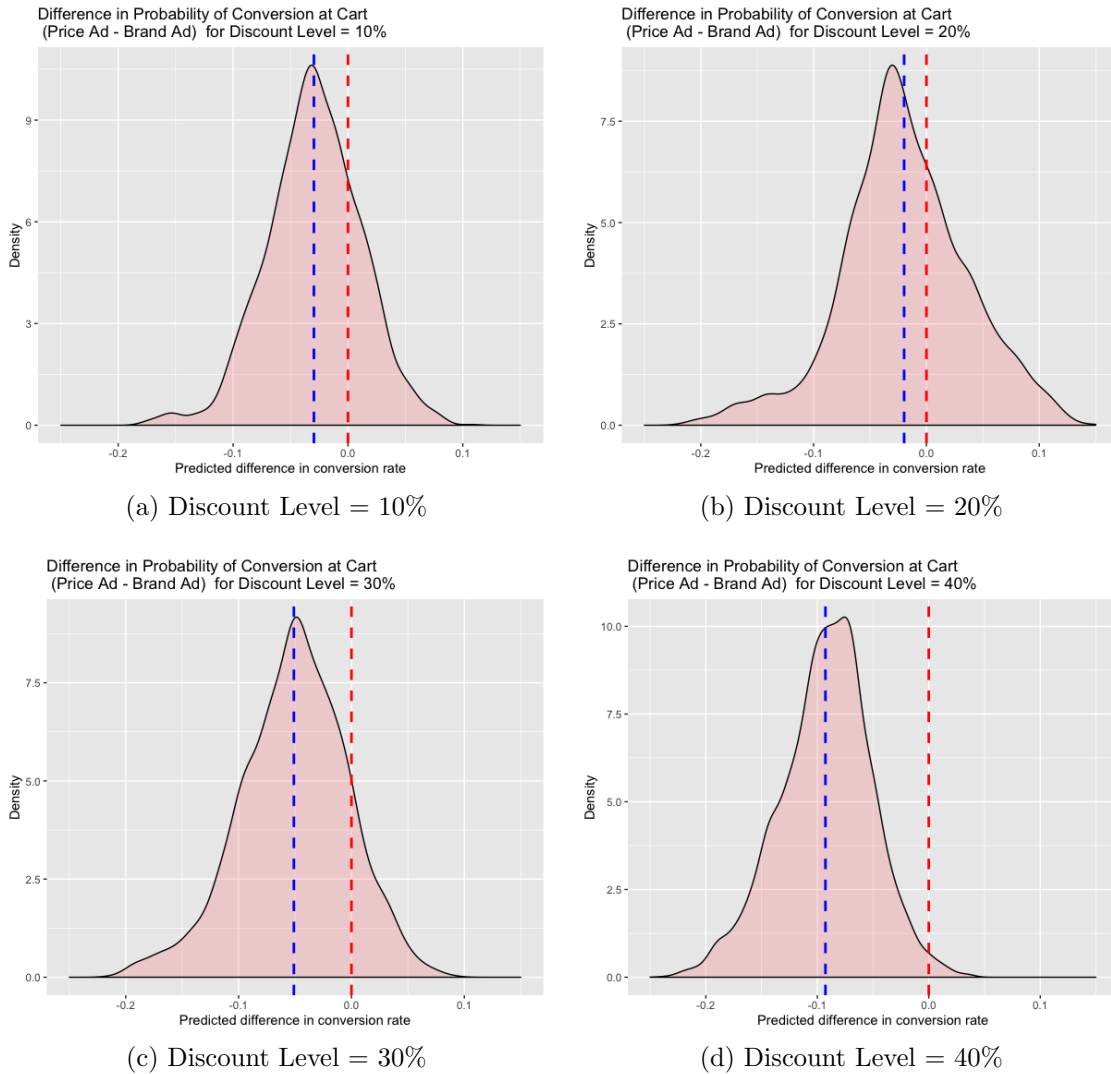


Figure 5.2: Distribution of differences in probability of conversion at cart after having been treated with a price ad and brand ad for different discount levels

cart the consumer discovers the full price to be paid and decides whether to convert based on this revealed price and the effect of the advertising that she has received in addition to her base brand preference.

5.2.1 Model specification

Individual i decides to visit the menu page on day t if

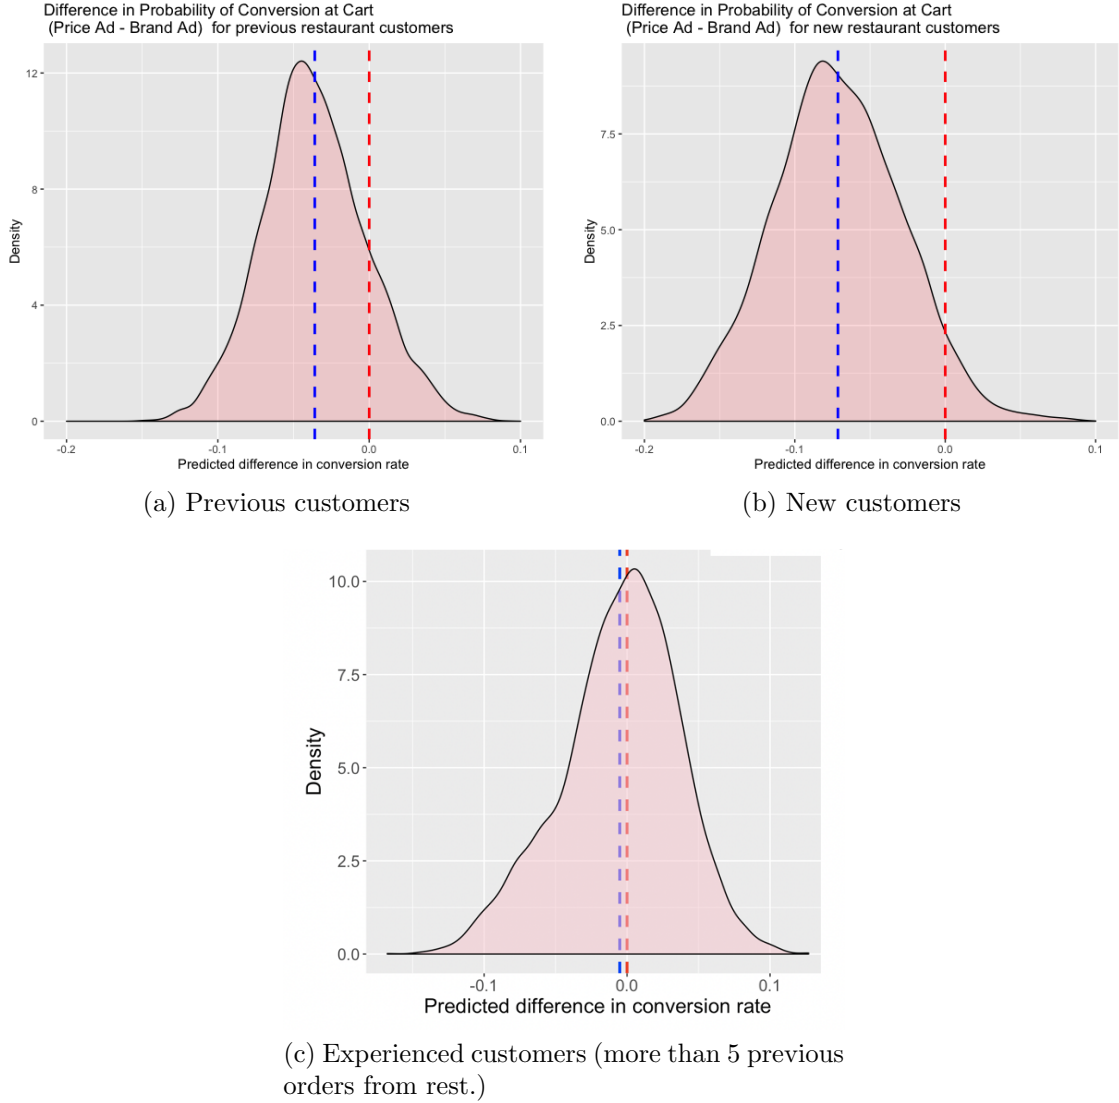


Figure 5.3: Distribution of differences in probability of conversion at cart after having been treated with a price ad and brand ad for previous, new and experienced focal restaurant customers

$$v_{1i} - c_{1i} - \beta_{1i}E_{it}[P|Ad^j] + \alpha_{1i}ad_{it}^j + \epsilon_{1it} > 0 \quad (5.1)$$

where

- v_{1i} captures individual i 's base preference for the focal restaurant
- c_{1i} indicates individual i 's search cost to visit the menu page

- $E_{it}[P|Ad^j]$ is the final price that individual i expects to pay under information contained in Ad^j where j indicates the type of ad received. The expectation of discount is zero before the first ad is received and then updates after the first ad is received, and stays at the updated level until the end of the experiment.
- ad_{it}^j is the individual's goodwill stock of ad type j at time t

Individual i decides to visit the cart page on day t if

$$v_{2i} - c_{2i} - \beta_{2i}E_{it}[P|Ad^j] + \alpha_{2i}ad_{it}^j + \epsilon_{2it} > 0 \quad (5.2)$$

where

- v_{2i} captures individual i 's base preference for the focal restaurant after having seen the menu
- c_{2i} indicates individual i 's search cost to visit the cart page from the menu page
- $E_{it}[P|Ad^j]$ is the final price that individual i expects to pay under information contained in Ad^j where j indicates the type of ad received
- ad_{it}^j is the individual's goodwill stock of ad type j at time t

Individual i decides to visit order on day t if

$$v_{3i} - c_{3i} - \beta_{3i}P + \alpha_{3i}ad_{it}^j + \epsilon_{3it} > 0 \quad (5.3)$$

- v_{3i} captures individual i 's base preference for the focal restaurant at the cart stage
- c_{3i} indicates individual i 's cost of finalizing the order by putting in her address, payment method etc.
- P is the actual price seen at cart

- ad_{it}^j is the individual's goodwill stock of ad type j at time t

We can set this model up to allow for selection on unobservables by letting the error terms in the three stages to be correlated. Specifically, they can be assumed to be draws from a trivariate normal distribution with zero means, unit variances and arbitrary covariance terms. We cannot identify search cost and base preference parameters separately, so they will be combined as one intercept term. Since I experimentally vary discount, and not the actual price, I replace the price terms above with discounts. I also control for observable characteristics by including a list of characteristics created using historical data. The set of observable characteristics is limited and not the full set used in the causal forest, in order to keep estimation tractable. Since there is no credible 'exclusion restriction' i.e. a variable that affects the search decision but not the purchase decision, I take a full information maximum likelihood approach to this instead of a Heckman selection approach. The estimation equations are:

$$y_{1it}^* = v_{1i} + \beta_{1i}E_{it}[D|Ad^j] + \alpha_{1i}ad_{it} + \delta_{1i}ad_{it}^{discount} + \gamma_1X_i + \epsilon_{1it} \quad (5.4)$$

$$y_{1it} = 1 \text{ if } y_{1it}^* > 0 \text{ and } 0 \text{ otherwise} \quad (5.5)$$

$$y_{2it}^* = v_{2i} + \beta_{2i}E_{it}[D|Ad^j] + \alpha_{2i}ad_{it} + \delta_{2i}ad_{it}^{discount} + \gamma_2X_i + \epsilon_{2it} \quad (5.6)$$

$$y_{2it} = 1 \text{ if } y_{2it}^* > 0 \text{ and } 0 \text{ otherwise} \quad (5.7)$$

$$y_{3it}^* = v_{3i} + \beta_{3i}D + \alpha_{3i}ad_{it} + \delta_{3i}ad_{it}^{discount} + \gamma_3X_i + \epsilon_{3it} \quad (5.8)$$

$$y_{3it} = 1 \text{ if } y_{3it}^* > 0 \text{ and } 0 \text{ otherwise} \quad (5.9)$$

where

- y_{1it} , y_{2it} and y_{3it} are dummy variables indicating whether individual i visited the menu page, visited the cart page or placed an order from the focal restaurant on day t
- X_i is the set of customer characteristics which include average order value, average discount percentage used in pre-experiment orders, fraction of pre-experiment orders purchased using discounts, dummy indicating whether the individual purchased atleast once from the focal restaurant pre-experiment and average restaurant cost for two among restaurants previously purchased from
- ad_{it} is the individual's goodwill stock of ads (of any type) at time t
- ad_{it}^{price} is the individual's goodwill stock of price ads at time t . The coefficient on this term gives the difference between price ads and brand ads. To get the full effect of price ad goodwill stock, we need to add the coefficients on ad_{it} and ad_{it}^{price}
- $E_{it}[D|Ad^j]$ is the expected discount under information contained in Ad^j where j indicates the type of ad received. The expectation of discount is zero before the first ad is received and then updates after the first ad is received, and stays at the updated level until the end of the experiment
- D is the actual discount seen at cart

5.2.2 Distributional Assumptions

Since very few individuals visit menu or cart more than once, it is not possible to estimate individual specific parameters. Instead we allow for individual heterogeneity by assuming that the intercept and the coefficients on discount and advertising come from a normal distribution with means and variances to be estimated. Off diagonal elements are fixed at zero. Suppose the vector $(v_{1i}, v_{2i}, v_{3i}, \beta_{1i}, \beta_{2i}, \beta_{3i}, \alpha_{1i}, \alpha_{2i}, \alpha_{3i})$ is denoted by θ , then $\theta \sim N(\mu, \Sigma)$. Mean parameters μ and diagonal elements of Σ are estimated. The error

terms $(\epsilon_{1i}, \epsilon_{2i}, \epsilon_{3i}) \sim N(0, \Sigma_\epsilon)$ i.e. the three equations are set up as a trivariate probit. The diagonal terms of Σ_ϵ are set to 1 and the off diagonal terms $\rho_{12}, \rho_{13}, \rho_{23}$ indicating the covariances between (ϵ_1, ϵ_2) , (ϵ_1, ϵ_3) and (ϵ_2, ϵ_3) are terms to be estimated.

5.2.3 Estimation Methodology

Estimation is done using simulated maximum likelihood. To simplify notation, suppose the equations 5.4, 5.6 and 5.8 are denoted as

$$y_{1it}^* = W'\eta + \epsilon_{1i} \quad (5.10)$$

$$y_{2it}^* = Y'\zeta + \epsilon_{2i} \quad (5.11)$$

$$y_{3it}^* = Z'\delta + \epsilon_{3i} \quad (5.12)$$

The likelihood (π_{it}) for an observation that has $y_{1i} = y_{2i} = y_{3i} = 1$ is $Pr(\epsilon_{1it} > -W'\eta, \epsilon_{2it} > -Y'\zeta, \epsilon_{3it} > -Z'\delta) = \Phi(W'\eta, Y'\zeta, Z'\delta; \rho_{12}, \rho_{23}, \rho_{13})$ where Φ denotes the standard normal CDF.

The likelihood (π_{it}) for an observation that has $y_{1i} = y_{2i} = 1$ and $y_{3i} = 0$ is $Pr(\epsilon_{1it} > -W'\eta, \epsilon_{2it} > -Y'\zeta, \epsilon_{3it} < -Z'\delta) = \Phi(W'\eta, Y'\zeta, -Z'\delta; \rho_{12}, -\rho_{23}, -\rho_{13})$

The likelihood (π_{it}) for an observation that has $y_{1i} = 1$ and $y_{2i} = 0$ is $Pr(\epsilon_{1it} > -W'\eta, \epsilon_{2it} < -Y'\zeta) = \Phi(W'\eta, -Y'\zeta; -\rho_{12})$

The likelihood (π_{it}) for an observation that has $y_{1i} = 0$ is $Pr(\epsilon_{1it} < -W'\eta) = 1 - \Phi(W'\eta)$

Since there are terms in η , ζ and δ that are random, specifically,

$(v_{1i}, v_{2i}, v_{3i}, \beta_{1i}, \beta_{2i}, \beta_{3i}, \alpha_{1i}, \alpha_{2i}, \alpha_{3i})$ denoted by $\theta \sim N(\mu, \Sigma)$, we will integrate the likelihood of each observation over the distribution of parameters $\pi_{it} = \int \pi_{it}(\theta) d\Omega(\theta)$ where $d\Omega(\theta)$ is the density of $\theta \sim N(\mu, \Sigma)$. This integration is implemented using a Monte Carlo simulation method; I take draws of θ from the distribution and compute an average value of

π_{it} for these draws.

The log likelihood for each observation is

$$\begin{aligned} \log(\pi_{it}) = & (y_{1it})(y_{2it})(y_{3it})\log(\Phi(W'\eta, Y'\zeta, Z'\delta; \rho_{12}, \rho_{23}, \rho_{13})) \\ & + (y_{1it})(y_{2it})(1 - y_{3it})\log(\Phi(W'\eta, Y'\zeta, -Z'\delta; \rho_{12}, -\rho_{23}, -\rho_{13})) \\ & + (y_{1it})(1 - y_{2it})\log(\Phi(W'\eta, -Y'\zeta; -\rho_{12})) + (1 - y_{1it})(1 - \Phi(W'\eta)) \end{aligned} \quad (5.13)$$

The total log likelihood is $\sum_{it} \log(\pi_{it})$ ¹

5.2.4 Results

The estimated mean and standard deviations for the intercept, price and ad terms at each stage are reported in Table 5.1.

Table 5.1: Three stage trivariate probit model estimation results

	Order Stage		Cart Stage		Menu Stage	
	Mean	Std. Dev	Mean	Std. Dev	Mean	Std. Dev
Intercept	0.029 (0.072)	0.009 (0.077)	-1.036*** (0.035)	0.135*** (0.027)	-2.217*** (0.011)	0.236*** (0.015)
<i>ad</i>	0.047(. (0.024)	0.012 (0.027)	-0.015 (0.013)	0.012 (0.016)	0.186*** (0.001)	0.026*** (0.002)
<i>ad</i> ^{discount}	-0.102*** (0.029)	0.047 (0.031)	-0.002 (0.017)	0.001 (0.022)	-0.003 (0.005)	0.001 (0.006)
<i>discount</i> (or $E_i[D Ad_j]$)	0.014*** (0.0008)	0.005*** (0.001)	0.006*** (0.0003)	0.002*** (0.0004)	0.002*** (0.0001)	0.001*** (0.0002)

Significance Codes: *** $p \leq 0.001$, ** $p \leq 0.01$, * $p \leq 0.05$, (.) $p \leq 0.1$

Estimates for coefficients on customer characteristics not shown here

We see from these results that even after allowing for selection on unobservables and controlling for observables, the estimated effect of price advertising on conditional conversion at cart is lower than that for brand advertising or no advertising. This is evidence of the

1. To save computation time, I subsampled individuals such that the individuals who had atleast one visit to cart and 10000 other randomly selected individuals in the experiment were in the estimation sample

‘discount spotlighting’ effect. Comparing the sizes of the estimates for the discount effect and the discount spotlighting effect on conversion to purchase, we see that the negative discount spotlighting effect would be offset by an approximately 7% additional discount (for an average customer) in order to achieve the same conversion rates to purchase at cart. This is similar to what we obtained with the causal forest.

I also re-estimate the model with an added reference term in the last step which captures the difference between the expectation of discount conditional on receiving a notification (the difference is zero for price ads). Results are shown in Table 5.2. They do not change and there seems to be no reference price effect.

Table 5.2: Three stage trivariate probit model estimation results with reference term added in the last stage

	Order Stage		Cart Stage		Menu Stage	
	Mean	Std. Dev	Mean	Std. Dev	Mean	Std. Dev
Intercept	0.025 (0.074)	0.008 (0.025)	-1.024*** (0.039)	0.132*** (0.023)	-2.121*** (0.012)	0.228*** (0.017)
<i>ad</i>	0.051(.) (0.025)	0.012 (0.028)	-0.013 (0.015)	0.011 (0.014)	0.169*** (0.001)	0.025*** (0.002)
<i>ad</i> ^{discount}	-0.095*** (0.031)	0.044 (0.029)	-0.005 (0.017)	0.002 (0.024)	-0.002 (0.005)	0.001 (0.006)
<i>discount</i> (or $E_i[D Ad_j]$)	0.015*** (0.001)	0.004*** (0.001)	0.007*** (0.0004)	0.003*** (0.0004)	0.002*** (0.0001)	0.001*** (0.0002)
Diff. bet. expected and actual discount	0.0005 (0.0006)	0.0003 (0.0003)				

Significance Codes: *** $p \leq 0.001$, ** $p \leq 0.01$, * $p \leq 0.05$, (.) $p \leq 0.1$

Estimates for coefficients on customer characteristics not shown here

We see that the effect of the reference term comes out to be statistically insignificant and the other results do not change. This shows that the reference price effect is not an important driving force in this experiment.

5.3 Ruling out alternate explanations

Apart from an effect on brand preference due to ‘discount spotlighting’, there could be two other reasons for price advertising leading to lower conversion to purchase relative to brand advertising. The first is a reference price effect. As discussed in earlier in this chapter, this has been ruled out by adding an additional term that controls for the difference between expected discount percentage before the cart stage and the actual discount percentage seen at cart. Also, since discount expectation under brand ads is between 20% and 30% (as explained in the previous chapter), we should see conversion under brand ads be lower than price ads at 10% discount level as the expected discount is higher than the actual discount. We can see from Figure 5.2, that conversion under brand ads is higher at 10% as well, indicating that the reference price effect is not a driving factor in this context.

Another explanation is that the consumers’ price sensitivity increases i.e. the co-efficient on price, conditional on arriving at cart is larger in absolute value under price advertising compared to brand advertising. If so, we should expect that conversion to purchase at high discounts should be high under price ads, i.e. even if brand ads are more effective, the difference between price ad and brand ad conversion rates would be lowest at high discount rates. Figure 5.2 shows the opposite pattern, thus ruling out this explanation. Note that lowered brand preference would lead to consumers being more price elastic even without a direct effect on the price co-efficient.

Finally, another explanation is that salience of the brand information in the ad is reduced in the price ad condition and the consumer doesn’t pay attention to it. If this is so, it should be reduced equally for all discount levels and there should be no heterogeneity in the effect with discount level.

CHAPTER 6

POST-EXPERIMENT DEMAND

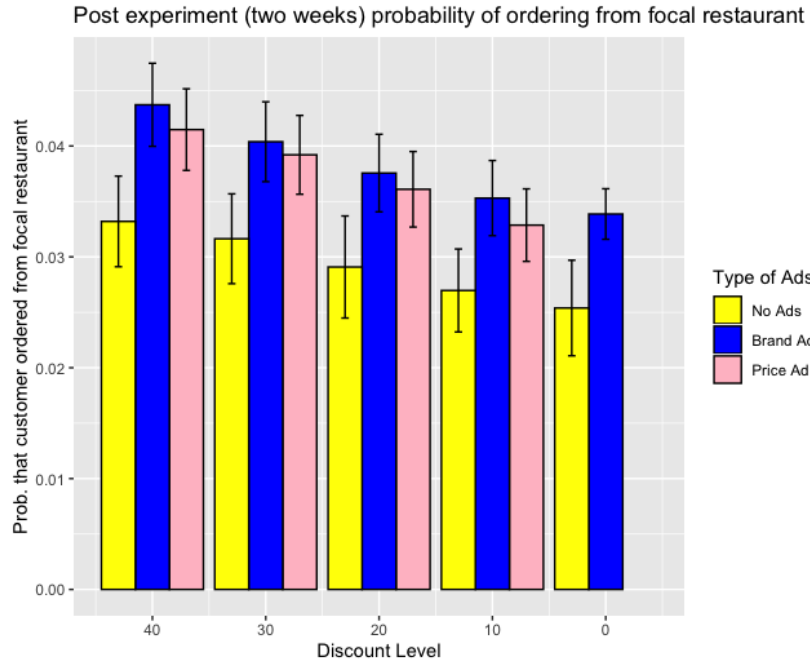
Next, we examine whether the different types of advertising have different impacts on post-experiment demand. Do the differences in demand arising from the different types of ads and discounts persist after the discounts and ads are no longer available or being sent to consumers?

There could be several different predictions: (1) Since ads and discounts make more customers buy during the experiment, these same consumers would also be more likely to buy post-experiment, as they have tried the product once and it is a popular restaurant (thus presumably of high quality). Also, there is no issue of stockpiling as this is a perishable product (2) Having received a discount during experiment, consumers would be less likely to buy if the discount is taken away as their reference price is lowered (Mazumdar et al., 2005), or if their brand preference is lowered due to the existence of discount (Dodson et al., 1978; Sawyer and Dickson, 1984) (3) The ‘discount spotlighting’ effect might make consumers who received price advertising less likely to buy due to lowered brand preference

Figure 6.1 shows the probability with which a customer that is assigned to each of these cells orders from the focal restaurant in the two weeks after the experiment’s end.

We see that the differences between ad types do not persist for two weeks post-experiment. At all discount levels, demand from consumers who were sent discount ads and those who were sent brand ads are statistically indistinguishable. This may be because there are two effects in opposite directions that approximately cancel each other out post-experiment. The positive effect of discount ads on search during the experiment leads the people who did search, to also be more likely to purchase post-experiment. Or, the negative effect on brand preference reduces probability of purchase post experiment.

We do see that there is a persistent positive ad effect and a positive discount effect. Those who saw any ad, and those who received a high discount are more likely to purchase in the post-promotion period, relative to those who received low discounts or did not receive any



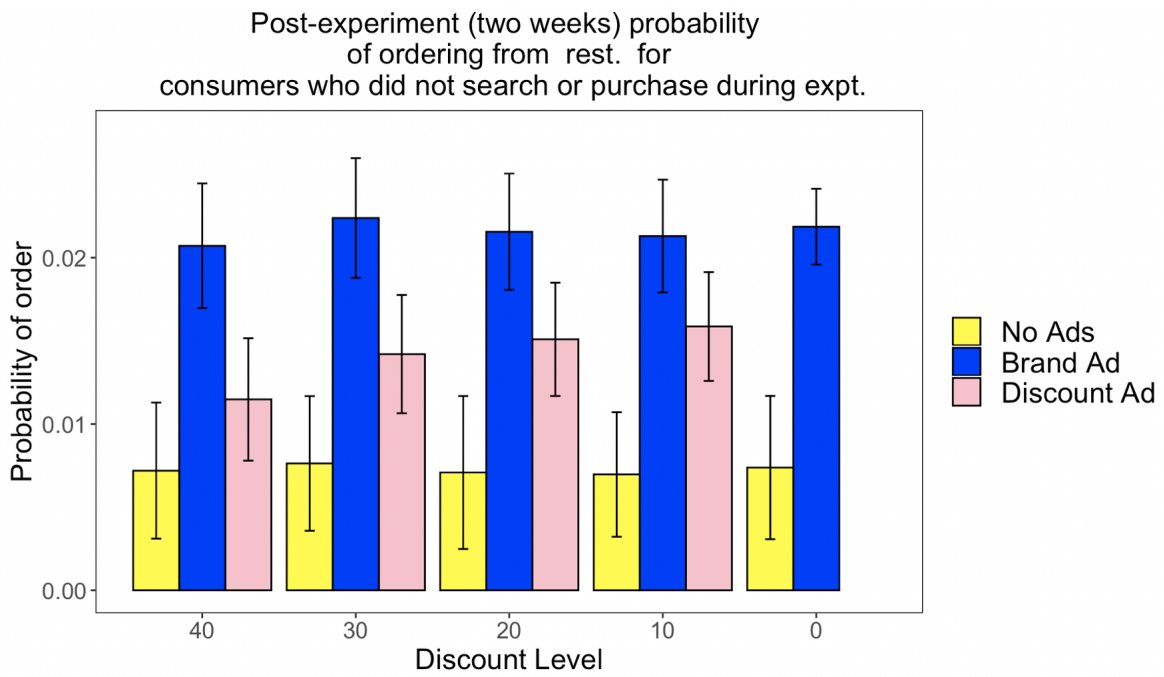
Error bars report 95% confidence intervals

Figure 6.1: Post experiment demand from focal restaurant

ads. This is inline with prediction (1) above. Increased demand during a promotion period through ads and discounts leads to greater demand in the post-promotion period as well.

Figure 6.2 shows that for consumers who did not search for the focal restaurant during the experiment, post experiment demand is in fact lower with discount ads.

This is consistent with the notion that the consumers who saw a discount ad, but did not search and find out more about the restaurant, think that it is a lower quality restaurant compared to those who received a brand ad (and also did not search).



Error bars report 95% confidence intervals

Figure 6.2: Post experiment demand from focal restaurant for consumers who did not search or purchase from it during expt.

CHAPTER 7

OPTIMAL TARGETING

In this experiment, the revenue maximizing un-targeted strategy is a 40% discount level with discount oriented advertising. The revenue curve for different ad types is shown in Figure 7.1.

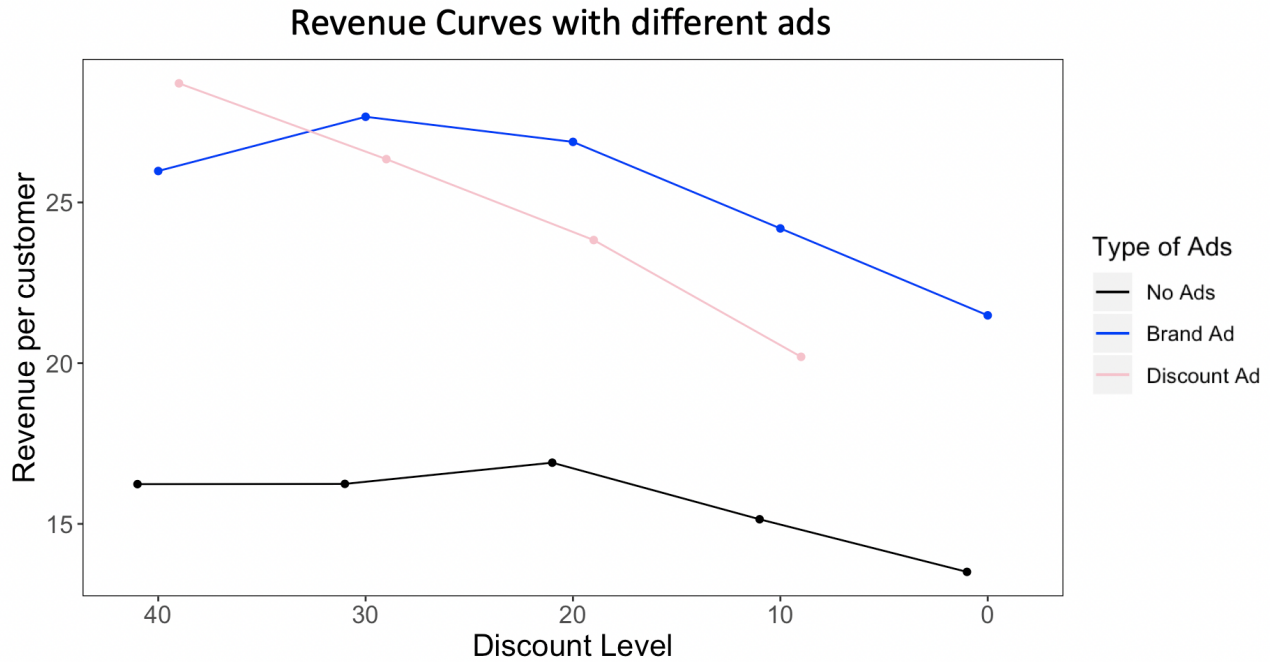


Figure 7.1: Revenue Curves with different advertising types

However, we have seen that discount oriented advertising results in increased search relative to brand oriented advertising (as long as the discount communicated is greater than expectations) and decreased conditional conversion to purchase (with the same discount level) relative to brand advertising. Further, these effects are heterogenous (for example, there is no reduced conversion to purchase for experienced consumers as they have a good idea of the quality of the restaurant already). Further the differences in probability of search depends on the expectations of discounts that consumers have based on previous ads received. The firm can take advantage of the positives of both types of advertising and optimally target both ad type and discount level to individual consumers.

The intuition behind the optimal targeting strategy is as follows:

- Use the causal forest to predict outcomes (search and conditional purchase) for each individual under each ad type and discount level
- If the consumer is likely to search with a brand ad, send it and then optimize discount amount for that individual based on predicted revenues for that consumers under different discount levels conditional on clicking on a brand ad
- If the consumer is likely to search only with discount ad, send a discount ad along with optimal discount level for that individual under discount advertising

In essence, the firm can use the causal forest to predict individual level demand curves and use the best possible combination of ad type and discount level for each individual to maximize revenue. A simulation of such a targeting strategy shows that expected revenue increases by 14% compared to optimal un-targeted strategy (discount ads along with 40% discount)

CHAPTER 8

MANAGERIAL IMPLICATIONS

The findings in this paper have several implications for managers who want to optimize pricing and advertising. First, at low discount levels, price advertising leads to lower overall demand relative to brand advertising. This directly contradicts general practice in the digital advertising industry of highlighting any available price discount in digital ads. In contexts characterized by high discounts, leading to high consumer expectations of discount, managers must be careful about highlighting discounts in ads that are lower than consumer expectations.

Further, the revenue maximizing price points under no advertising, price advertising and brand advertising were different, price and type of advertising must be jointly optimized. I show that optimally targeting both ad type and discount level for individual consumers based on their predicted probabilities of search and purchase increases revenue compared to the optimal un-targeted ad type and discount level.

The self-selection of different types and numbers of consumers into searching for the product under brand and price advertising at different discount levels has implications for targeting of advertising content. At low discount levels, brand advertising can lead to higher search and demand compared to price advertising. i.e. it can increase the probability of search if targeted towards people with high expectations of discount conditional on receiving a notification. Of course, these beliefs will get updated as well, so care must be taken not to disappoint the consumer too often by repeatedly revealing to them a price at cart that is higher than their expectations.

The ‘discount spotlighting effect’, which leads to lower conversion to purchase at cart, has implications for targeting of point-of-sale discounts. For example, firms often use cart abandonment reminders or discounts to induce customers who have abandoned their cart to finish a purchase. The fact that people who see different ads convert at different rates at the cart stage implies that cart abandonment discounts can be targeted according to the type of

advertising that the consumer received. A given individual who has seen a brand ad is likely to need a lower cart abandonment discount than one who has seen a price ad (assuming that such a discount is profitable for the firm).

The above implications are directly relevant for the partner firm. With its rich set of historical consumer data, and the findings from this experiment, the firm predict the probability of search and purchase for each customer (i.e. a set of customer characteristics) under different discount levels and advertising type. Thus, it can jointly optimize prices along with advertising type and target both discounts and advertising content. Since it already sends cart abandonment notifications, it can further include a personalized discount along with these according to customer characteristics and the type of advertising seen.

Managers might also wish to screen customers while running a promotional campaign. If servicing the search process for a customer is expensive (perhaps in an offline setting where servicing costs are high), the firm would rather have a lower number of customers who are more likely to convert make the initial search rather than a higher number of people who are less likely to convert. If the available discount is high, they can use brand advertising to attract a smaller group of individuals who are more likely to convert after they discover the true price. For example at a 30% discount in the experiment, the number of orders is similar, but the number of menu page visits made by consumers who received a price ad is higher.

Firms and academics often use two price points to estimate the demand curve by making a linearity assumption. Since brand ads lead to a similar number of people visiting the cart at all discount levels, there is a point of demand inelasticity when conversion hits 100%. This point leads to a non linearity in the demand curve under brand advertising, which may lead to mis-measurement of the slope of demand under the linearity assumption. This has to be kept in mind while designing experiments to measure price elasticities under different forms of advertising.

CHAPTER 9

CONCLUSION

In this dissertation, I show that the inclusion of discount information in advertising has a causal effect of decreasing brand preference through the ‘discount spotlighting effect’. I also justify the prevalence of price advertising by showing that the negative effect of ‘discount spotlighting’ on conversion to purchase is offset by attracting many more consumers to search for the product at high discount levels. Further, I illustrate the effects of price and brand advertising on overall demand, search and purchase conditional on search. I show that discount advertising may lower search and purchase at low discount levels, thus challenging conventional wisdom in digital advertising that any available discount should be highlighted. Finally, I find that price advertising causes consumers to shift their purchases among rival (non-advertised) firms to those that offer discounts. This has implications for price competition among firms under different advertising regimes and optimal platform policy regarding advertising content.

A key limitation of this study is that the variation in advertising types and prices is purely cross-sectional. Thus, I am unable to make comments on how different sequences of advertising types or different sequences of discounts affect consumer behavior. Another limitation is that there was only one advertised brand. Thus, effects of different firms competing with different types of advertising were not studied. These are interesting topics to explore in future research.

Other promising avenues for future research include looking at how the ‘discount spotlighting effect’ is moderated by different types of messaging. For example, does the size of the effect reduce if the price discount is explicitly coming from the platform, not the restaurant? eg: ‘Uber Eats is offering everyone 10% off all purchases from restaurant X today’. What if the price discount is targeted to a specific customer? eg: ‘To thank you for your loyalty, we’re offering you 10% off’. Or, what if the price discount is explained due to some external factors such as a holiday? eg: ‘There is a 10% discount this weekend due to the Presidents

Day holiday'.

REFERENCES

- Aaker, D. A. (2009). *Managing brand equity*. simon and schuster.
- Bertrand, M., Karlan, D., Mullainathan, S., Shafir, E., and Zinman, J. (2010). What's advertising content worth? evidence from a consumer credit marketing field experiment. *The Quarterly Journal of Economics*, 125(1):263–306.
- Blattberg, R. C. and Neslin, S. A. (1989). Sales promotion: The long and the short of it. *Marketing letters*, 1(1):81–97.
- Comanor, W. S. and Wilson, T. A. (1979). The effect of advertising on competition: A survey. *Journal of Economic Literature*, 17(2):453–476.
- Dodson, J. A., Tybout, A. M., and Sternthal, B. (1978). Impact of deals and deal retraction on brand switching. *Journal of Marketing Research*, 15(1):72–81.
- Erdem, T., Keane, M. P., and Sun, B. (2008). The impact of advertising on consumer price sensitivity in experience goods markets. *Quantitative Marketing and Economics*, 6(2):139–176.
- Greene, W. H. (2000). Econometric analysis 4th edition. *International edition, New Jersey: Prentice Hall*, pages 201–215.
- Heckman, J. J. (1979). Sample selection bias as a specification error. *Econometrica: Journal of the Econometric Society*, pages 153–161.
- Imbens, G. W. and Rubin, D. B. (2015). *Causal inference in statistics, social, and biomedical sciences*. Cambridge University Press.
- Jedidi, K., Mela, C. F., and Gupta, S. (1999). Managing advertising and promotion for long-run profitability. *Marketing science*, 18(1):1–22.
- Kaul, A. and Wittink, D. R. (1995). Empirical generalizations about the impact of advertising on price sensitivity and price. *Marketing Science*, 14(3_supplement):G151–G160.
- Mayzlin, D. and Shin, J. (2011). Uninformative advertising as an invitation to search. *Marketing Science*, 30(4):666–685.
- Mazumdar, T., Raj, S. P., and Sinha, I. (2005). Reference price research: Review and propositions. *Journal of Marketing*, 69(4):84–102.
- Mela, C. F., Gupta, S., and Lehmann, D. R. (1997). The long-term impact of promotion and advertising on consumer brand choice. *Journal of Marketing Research*, 34(2):248–261.
- Nelson, P. (1970). Information and consumer behavior. *Journal of Political Economy*, 78(2):311–329.
- Olson, J. C. (1977). Price as an informational cue: effects on product evaluations,(w:) woodside ag, sheth jn, bennet pd (eds), consumer and industrial buying behavior.

- Ozga, S. A. (1960). Imperfect markets through lack of knowledge. *The Quarterly Journal of Economics*, pages 29–52.
- Rao, A. R. and Monroe, K. B. (1989). The effect of price, brand name, and store name on buyers? perceptions of product quality: An integrative review. *Journal of marketing Research*, 26(3):351–357.
- Rubin, D. B. (1990). Formal mode of statistical inference for causal effects. *Journal of Statistical Planning and Inference*, 25(3):279–292.
- Sawyer, A. G. and Dickson, P. R. (1984). Psychological perspectives on consumer response to sales promotion. *Research on Sales Promotion: Collected papers*, 2.
- Scitovszky, T. (1945). Some consequences of the habit of judging quality by price. *The Review of Economic Studies*, 12(2):100–105.
- Shapiro, B., Hitsch, G. J., and Tuchman, A. (2018). Generalizable and robust tv ad effects. *Available at SSRN*.
- Stigler, G. J. (1961). The economics of information. *Journal of Political Economy*, 69(3):213–225.
- Sudhir, K., Roy, S., and Cherian, M. (2016). Do sympathy biases induce charitable giving? the effects of advertising content. *Marketing Science*, 35(6):849–869.
- Wager, S. and Athey, S. (2018). Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association*, 113(523):1228–1242.

APPENDIX A

RANDOMIZATION CHECKS

I ensure the assignment of discount level, ad frequency and ad type is sufficiently randomized by examining the correlations between the treatment variables and pre-treatment customer characteristics

Table A.1: Randomization Checks 1

	Brand Ad assigned		Discount Ad assigned	
	Correlation	p-value	Correlation	p-value
Avg. Order Value	-0.0003	0.87	-0.0004	0.83
Previous focal rest. customer	-0.0002	0.92	0.003	0.17
Avg. discount % used	0.0001	0.94	-0.0004	0.82
Avg. rest. cost for two	0.001	0.41	0.0008	0.7
Fraction of orders made on discount	-0.0008	0.71	-0.0003	0.88
Expected discount conditional on receiving a notification	0.001	0.47	-0.001	0.65

Table A.2: Randomization Checks 2

	Discount Level		Ad Frequency	
	Correlation	p-value	Correlation	p-value
Avg. Order Value	0.0000	0.98	-0.002	0.33
Previous focal rest. customer	0.002	0.26	-0.001	0.6
Avg. discount % used	0.0009	0.68	-0.002	0.22
Avg. rest. cost for two	-0.001	0.41	0.002	0.29
Fraction of orders made on discount	-0.0003	0.86	-0.003	0.12
Expected discount conditional on receiving a notification	-0.001	0.53	0.003	0.14